Best practices for deploying Citrix XenServer on HP StorageWorks P4000 SAN

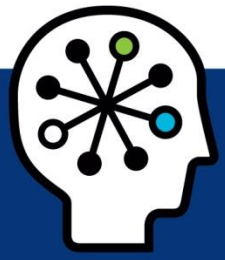

Table of contents

Executive summary 3

Business case 3
- High availability 4
- Scalability 4
- Virtualization 4

XenServer storage model 5
- HP StorageWorks P4000 SAN 5
- Storage repository 6
- Virtual disk image 6
- Physical block device 6
- Virtual block device 6

Overview of XenServer iSCSI storage repositories 6
- iSCSI using the software initiator (lvmoiscsi) 6
- iSCSI Host Bus Adapter (HBA) (lvmohba) 6
- SAN connectivity 6
- Benefits of shared storage 7
- Storage node 7
  - Clustering and Network RAID 8
  - Networking bonding 8

Configuring an iSCSI volume 9
- Example 9
  - Creating a new volume 10
  - Configuring the new volume 11
  - Comparing full and thin provisioning 12
  - Benefits of thin provisioning 12

Configuring a XenServer Host 13
- Synchronizing time 14
  - NTP for XenServer 15
- Network configuration and bonding 16
  - Example 17
- Connecting to an iSCSI volume 19
  - Determining or changing the host's IQN 19
  - Specifying IQN authentication 21
  - Creating an SR 25
- Creating a VM on the new SR 28
- Summary 30

Configuring for high availability 31

Configuration ... 33
Implementing Network RAID for SRs ... 33
Configuring Network RAID ... 34
Pooling XenServer hosts ... 35
Configuring VMs for high availability ... 36
Creating a heartbeat volume ... 36
Configuring the resource pool for HA ... 37
Configuring multi-site high availability with a single cluster ... 39
Configuring multi-site high availability with multiple clusters ... 41
Disaster recoverability ... 45
Backing up configurations ... 46
Resource pool configuration ... 46
Host configuration ... 46
Backing up metadata ... 47
SAN based Snapshots ... 48
SAN based Snapshot Rollback ... 49
Reattach storage repositories ... 50
Virtual Machines (VMs) ... 51
Creating VMs ... 51
Size of the storage repository ... 51
Increasing storage repository volume size ... 52
Uniqueness of VMs ... 53
Process preparing a VM for Cloning ... 54
Changing the Storage Repository and Virtual Disk UUID ... 54
SmartClone the Golden Image VM ... 60
For more information ... 63

Executive summary

Using Citrix XenServer with HP StorageWorks P4000 SAN storage, you can host individual desktops and servers inside virtual machines (VMs) that are hosted and managed from a central location utilizing optimized, shared storage. This solution provides cost-effective high availability and scalable performance.

Organizations are demanding better resource utilization, higher availability, along with more flexibility to react to rapidly changing business needs. The 64-bit XenServer hypervisor provides outstanding support for VMs, including granular control over processor, network, and disk resource allocations; as a result, your virtualized servers can operate at performance levels rates that closely match physical platforms. Meanwhile, additional XenServer hosts deployed in a resource pool provide scalability and support for high-availability (HA) applications, allowing VMs to restart automatically on other hosts at the same site or even at a remote site.

Enterprise IT infrastructures are powered by storage. HP StorageWorks P4000 SANs offer scalable storage solutions that can simplify management, reduce operational costs, and optimize performance in your environment. Easy to deploy and maintain, HP StorageWorks P4000 SAN storages help to ensure that crucial business data remains available; through innovative double-fault protection across the entire SAN, your storage is protected from disk, network, and storage node faults.

You can grow your HP StorageWorks P4000 SAN non-disruptively, in a single operation, by simply adding storage; thus, you can scale performance, capacity and redundancy as storage requirements evolve. Features such as asynchronous and synchronous replication, storage clustering, network RAID, thin provisioning, snapshots, remote copies, cloning, performance monitoring, and a single pane-of-glass management can add value in your environment.

This paper explores options for configuring and using XenServer, with emphasis on best practices and tips for an HP StorageWorks P4000 SAN environment.

Target audience: This paper provides information for XenServer Administrators interested in implementing XenServer-based server virtualization using HP StorageWorks P4000 SAN storage. Basic knowledge of XenServer technologies is assumed.

Business case

Organizations implementing server virtualization typically require shared storage to take full advantage of today's powerful hypervisors. For example, XenServer supports features such as XenMotion and HA that require shared storage to serve a pool of XenServer host systems. By leveraging the iSCSI storage protocol, XenServers are able to access the storage just like local storage but over an Ethernet network. Since standard Ethernet networks are already used by most IT organizations to provide their communications backbones, no additional specialized hardware is required to support a Storage Area Network (SAN) implementation. Security of your data is handled first-most by authentication at the storage, physical, as well as the iSCSI protocol mechanisms itself. Just like any other data, it too can be encrypted at the client, thereby satisfying any data security compliance requirements.

Rapid deployment

Shared storage is not only a requirement for a highly-available XenServer configuration; it is also desirable for supporting rapid data deployment. Using simple management software, you can respond to a request for an additional VM and associated storage with just a few clicks. To minimize deployment time, you can use a golden-image clone, with both storage and operating system (OS) pre-configured and ready for application deployment.

"Data de-duplication[1]" allows you to roll out hundreds of OS images while only occupying the space needed to store the original image. Initial deployment time is reduced to the time required to perform the following activities:

- Configure "the first operating system"
- Configure the particular deployment for uniqueness
- Configure the applications in VMs

No longer should a server roll-out take days.

High availability

Highly-available storage is a critical component of a highly-available XenServer resource pool. If a XenServer host at a particular site were to fail or the entire site were to go down, the ability of another XenServer pool to take up the load of the affected VMs means that your business-critical applications can continue to run.

HP StorageWorks P4000 SAN solutions provide the following mechanisms for maximizing availability:

- Storage nodes are clustered to provide redundancy.
- Hardware RAID implemented at the storage-node level can eliminate the impact of disk drive failures.
- Configuring multiple network connections to each node can eliminate the impact of link failures.
- Synchronous replication between sites can minimize the impact of a site failure.
- Snapshots can prevent data corruption when you are rolling back to a particular point-in–time.
- Remote snapshots can be used to add sources for data recovery.

Comprehensive, cost-effective capabilities for high availability and disaster recovery (DR) applications are built into every HP StorageWorks P4000 SAN. There is no need for additional upgrades; simply install a storage node and start using it. When you need additional storage, higher performance, or increased availability, just add one or more storage nodes to your existing SAN.

Scalability

The storage node is the building block of an HP StorageWorks P4000 SAN, providing disk spindles, a RAID backplane, CPU processing power, memory cache, and networking throughput that, in combination, contribute toward overall SAN performance. Thus, HP StorageWorks P4000 SANs can scale linearly and predictably as your storage requirements increase.

Virtualization

Server virtualization allows you to consolidate multiple applications using a single host server or server pool. Meanwhile, storage virtualization allows you to consolidate your data using multiple storage nodes to enhance resource utilization, availability, performance, scalability and disaster recoverability, while helping to achieve the same objectives for VMs.

[1] Occurring at the iSCSI block level on the host server disk

The following section outlines the XenServer storage model.

XenServer storage model

The XenServer storage model used in conjunction with HP StorageWorks P4000 SANs is shown in Figure 1.

Figure 1. XenServer storage model

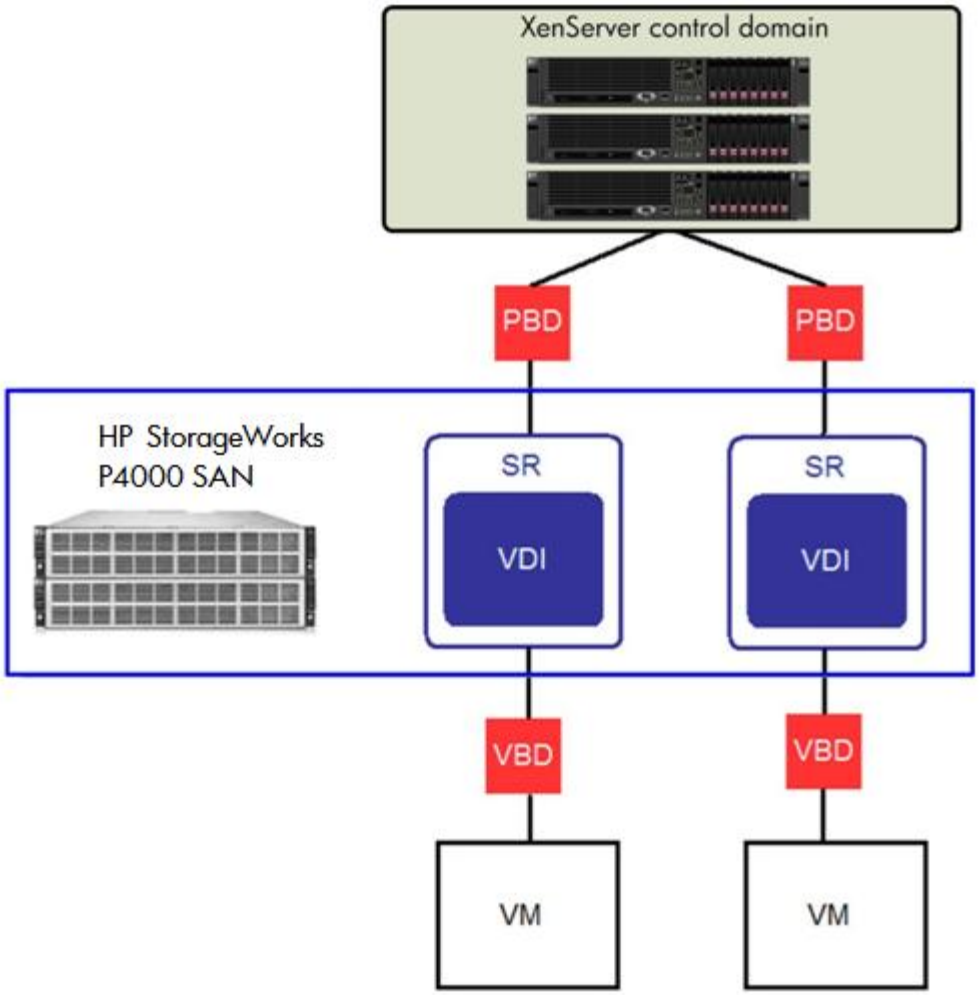

Brief descriptions of the components of this storage model are provided below.

HP StorageWorks P4000 SAN

A SAN can be defined as an architecture that allows remote storage devices to appear to a server as though these devices are locally-attached.

In an HP StorageWorks P4000 SAN implementation, data storage is consolidated on a pooled cluster of storage nodes to enhance availability, resource utilization, and scalability. Volumes are allocated to XenServer hosts via an Ethernet infrastructure (1 Gb/second or 10 Gb/second) that utilizes the iSCSI block-based storage protocol.

Performance, capacity and availability can be scaled on-demand and on-line.

Storage repository

A storage repository (SR) is defined as a container of storage to which XenServer Virtual Machine data will be stored. Although SRs can support locally connected storage types such as IDE, SATA, SCSI and SAS drives, remotely connected iSCSI SAN storage will be discussed in this document. Storage Repositories abstract the underlying differences in the storage connectivity, although differences between local and remotely connected storage repositories enable specific XenServer Resource Pool capabilities such as High Availability and XenMotion. The demands of a XenServer resource pool dictate that storage must be equally accessible amongst all hosts and therefore data must not be stored on local SRs

Virtual disk image

A virtual disk image (VDI) is the disk presented to the Virtual Machine and its OS as a local disk (even if the disk image is stored remotely). This image will be stored in the container of a SR. Although multiple VDIs may be stored on a single SR, it is considered best practice to store one VDI per SR. It is also considered best practice to have one Virtual Machine allocated to a single VDI. VDIs can be stored in different formats depending upon the type of connectivity afforded to Storage Repository.

Physical block device

A physical block device (PBD) is a connector that describes how XenServer hosts find and connect to an SR.

Virtual block device

A virtual block device (VBD) is a connector that describes how a VM connects to its associated VDI, which is located on a SR.

Overview of XenServer iSCSI storage repositories

XenServer hosts access HP StorageWorks P4000 iSCSI storage repositories (SRs) either using the open-iSCSI software initiator or thru an iSCSI Host Bus Adapter (HBA). XenServer SRs regardless of the access method, use Linux Logical Volume Manager (LVM) as the underlying file system to store Virtual Disk Images (VDI). Although multiple virtual disks may reside on a single storage repository, it is recommended that a single VDI occupy the space of an SR for optimum performance.

iSCSI using the software initiator (lvmoiscsi)

In this method, the open-iSCSI software initiator is used to connect to an iSCSI volume over the Ethernet network. The iSCSI volume is presented to a XenServer or Resource Pool thru this connection.

iSCSI Host Bus Adapter (HBA) (lvmohba)

In this method, a specialized hardware interface, an iSCSI Host Bus Adapter, is used to connect to an iSCSI volume over the Ethernet network. The iSCSI volume is presented to a XenServer or Resource Pool through this connection.

lvmoiscsi is the method used in this paper. Please refer to the XenServer Administrator's Guide for configuration requirements of lvmohba.

SAN connectivity

Physically connected via an Ethernet IP infrastructure, HP StorageWorks P4000 SANs provide storage for XenServer hosts using the iSCSI block-based storage protocol to carry storage data from host to

storage or from storage to storage. Each host acts as an initiator (iSCSI client) connecting to a storage target (HP StorageWorks P4000 SAN volume) in a SR, where the data is stored.

Since SCSI commands are encapsulated within an Ethernet packet, storage no longer needs to be locally-connected, inside a server. Thus, storage performance for a XenServer host becomes a function of bandwidth, based on 1 Gb/second or 10 Gb/second Ethernet connectivity.

Moving storage from physical servers allows you to create a SAN where servers must now remotely access shared storage. The mechanism for accessing this shared storage is iSCSI, in much the same way as other block-based storage protocols such as Fibre Channel (FC). SAN topology can be deployed efficiently using the standard, pre-existing Ethernet switching infrastructure.

Benefits of shared storage

The benefits of sharing storage include:

- The ability to provide equal access to shared storage is a basic requirement for hosts deployed in a resource pool, enabling XenServer functionality such as HA and XenMotion, which supports the migration of VMs between resource pools in the event of a failover or a manual live state.
- Since storage resources are no longer dedicated to a particular physical server, utilization is enhanced; moreover, you are now able to consolidate data.
- Storage reallocation can be achieved without cabling changes.

In much the same way that XenServer can be used to efficiently virtualize server resources, HP StorageWorks P4000 SANs can be used to virtualize and consolidate storage resources while extending storage functionality. Backup and DR are also simplified and enhanced by the ability to move VM data anywhere an Ethernet packet can travel.

Storage node

The storage node is the basic building block of an HP StorageWorks P4000 SAN and includes the following components:

- CPU
- Disk drives
- RAID controller
- Memory
- Cache
- Multiple network interfaces

These components work in concert to respond to storage read and write requests from an iSCSI client.

The RAID controller supports a range of RAID types for the node's disk drives, allowing you to configure different levels of fault-tolerance and performance within the node. For example, RAID 10 maximizes throughput and redundancy, RAID 6 can compensate for dual disk drive faults while better utilizing capacity, and RAID 5 provides minimal redundancy but maximizes capacity utilization.

Network interfaces can be used to provide fault tolerance or may be aggregated to provide additional bandwidth. 1 Gb/second and 10 Gb/second interfaces are supported.

CPU, memory, and cache work together to respond to iSCSI requests for reading or writing data.

All physical storage node components described above are virtualized, becoming a building block for an HP StorageWorks P4000 SAN.

Clustering and Network RAID

Since an individual storage node would represent a single point of failure (SPOF), the HP StorageWorks P4000 SAN supports a cluster of storage nodes working together and managed as a single unit. Just as conventional RAID can protect against a SPOF within a disk, Network RAID can be used to spread a volume's data blocks across the cluster to protect against single or multiple storage node failures.

HP StorageWorks SAN/iQ, the storage software logic, performs the storage virtualization and distributes data across the cluster.

Network RAID helps prevent storage downtime for XenServer hosts accessing that volume, which is critical for ensuring that these hosts can always access VM data. An additional benefit of virtualizing a volume across the cluster is that the resources of all the nodes can be combined to increase read and write throughput as well as capacity[2].

It is a best practice to configure a minimum of two nodes in a cluster and use Network RAID at Replication Level 2.

Networking bonding

Each storage node supports multiple network interfaces to help eliminate SPOFs from the communication pathway. Configuration of the network interfaces is best implemented when attaching both network interfaces to an Ethernet switching infrastructure.

Network bonding[3] provides a mechanism for aggregating multiple network interfaces into a single, logical interface. Bonding supports path failover in the event of a failure; in addition, depending on the particular options configured, bonding can also enhance throughput.

In its basic form, a network bond forms an active/passive failover configuration; that is, if one path in this configuration were to fail, the other would assume responsibility for communicating data.

Note
Each network interface should be connected to a different switch.

With Adaptive Load Balancing (ALB) enabled on the network bond, both network interfaces can transmit data from the storage node; however, only one interface can receive data. This configuration requires no additional switch configuration support and may also span each connection across multiple switches ensuring there is no single point of failure to multiple switches.

Enabling IEEE 802.3ad Link Aggregation Control Protocol (LACP) Dynamic Mode on the network bond allows both network ports to send and receive data in addition to providing fault tolerance. However, the associated switch must support this feature; pre-configuration may be required for the attached ports.

Note
LACP requires both network interfaces' ports to be connected to a single switch, thus creating a potential SPOF.

Best practices for network configuration depend on your particular environment; however, at a minimum, you should configure an ALB bond between network interfaces.

[2] The total space available for data storage is the sum of storage node capacities.
[3] Also known as NIC teaming, where "NIC" refers to a network interface card

Configuring an iSCSI volume

The XenServer SR stores VM data on a volume (iSCSI Target) that is a logical entity with specific attributes. The volume consists of storage on one or more storage nodes.

When planning storage for your VMs, you should consider the following:

- How will that storage be used?
- What are the storage requirements at the OS and application levels?
- How would data growth impact capacity and data availability?
- Which XenServer host – or, in a resource pool, hosts – require access to the data?
- How does your DR approach affect the data?

Example

An HP StorageWorks P4000 SAN is configured using the Centralized Management Console (CMC). In this example, the **HP-Boulder** management group defines a single storage site for a XenServer host resource pool (farm) or a synchronously-replicated stretch resource pool. **HP-Boulder** can be thought of as a logical grouping of resources.

A cluster named **IT-DataCenter** contains two storage nodes, **v8.1-01** and **v8.1-02**.

20 volumes have currently been created. This example focuses on volume **XPSP2-01**, which is sized at 10GB; however, because it has been thinly provisioned, this volume occupies far less space on the SAN. Its iSCSI qualified name (IQN) is **iqn.2003-10.com.lefthandnetworks:hp-boulder:55:xpsp2-01**, which uniquely identifies this volume in the SR.

Figure 2 shows how the CMC can be used to obtain detailed information about a particular storage volume.

Figure 2. Using CMC to obtain detailed information about volume **XPSP2-01**

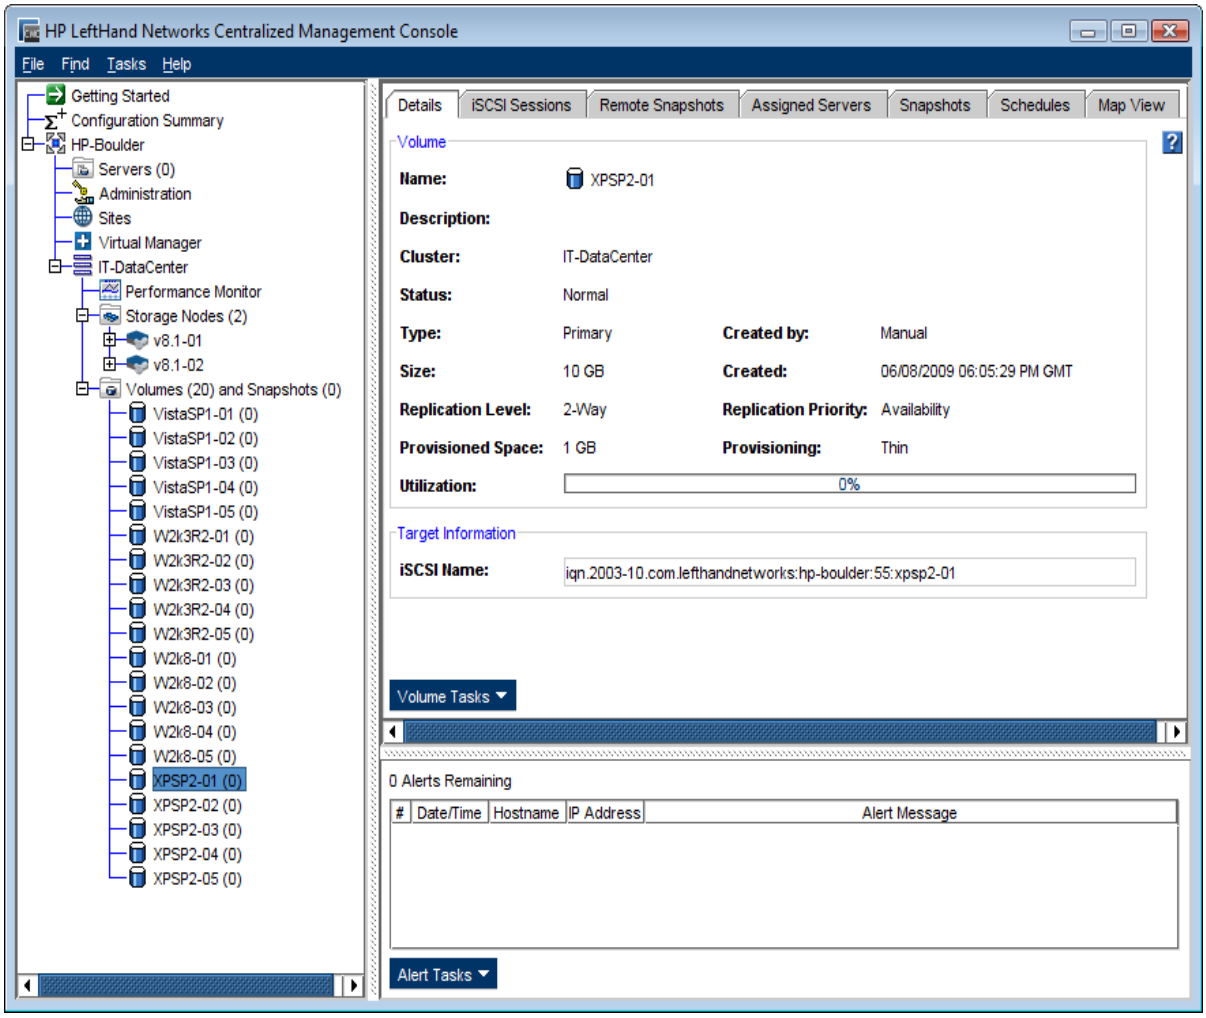

Creating a new volume

The CMC is used to create volumes such as **XPSP2-01,** as shown in Figure 3**.**

Figure 3. Creating a new volume

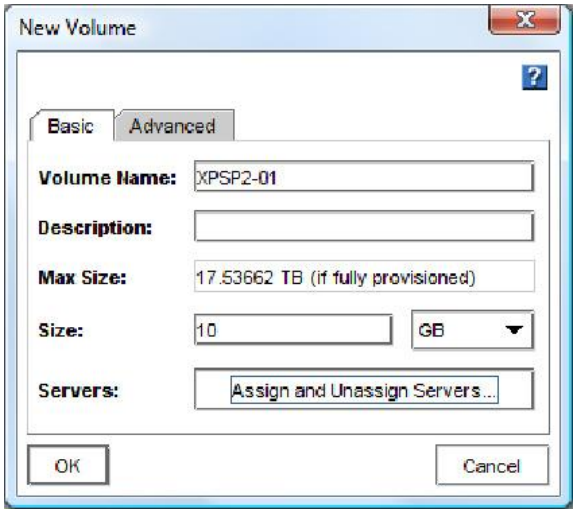

It is a best practice to create a unique iSCSI volume for each VM in an SR. Thus, HP suggests matching the name of the VM to that of the XenServer SR and of the volume created in the CMC. Using this convention, it is always clear which VM is related to which storage allocation.

This example is based on a 10GB Windows XP SP2 VM. The name of the iSCSI volume – **XPSP2-01** – is repeated when creating the SR as well as the VM.

The assignment of Servers will define which iSCSI Initiators (XenServer Hosts) are allowed to read/write to the storage and will be discussed later in the *Configuring a XenServer Host* section.

Configuring the new volume

Network RAID (**2-Way** replication) is selected to enhance storage availability; now, the cluster can survive at most one non-adjacent node failure.

Note
The more nodes there are in a cluster, the more nodes can fail without XenServer hosts losing access to data.

Thin Provisioning has also been selected to maximize data efficiency in the SAN–only data that is actually written to the volume that can occupy space. In functionality, this is equivalent to a sparse XenServer virtual hard drive (VHD); however, it is implemented efficiently in the storage with no limitation on the type of volume connected within XenServer.

Figure 4 shows how to configure **Thin Provisioning**.

Figure 4. Configuring **2-Way Replication** and **Thin Provisioning**

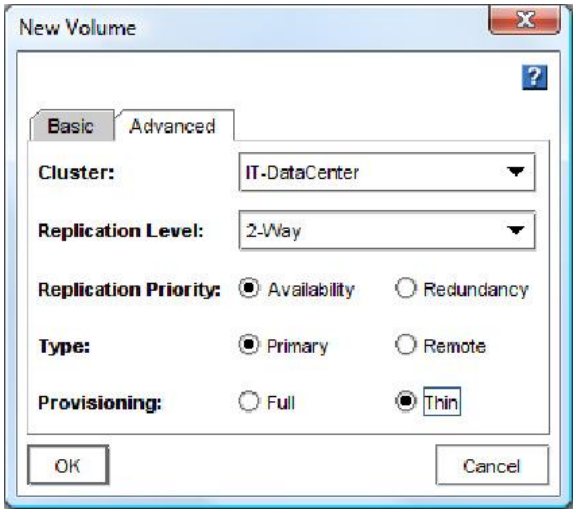

You can change volume properties at any time. However, if you change volume size, you may also need to update the XenServer configuration as well as the VM's OS in order for the new size to be recognized.

Comparing full and thin provisioning

You have two options for provisioning volumes on the SAN:

- **Full Provisioning**
  With **Full Provisioning**, you reserve the same amount of space in the storage cluster as that presented to the XenServer host. Thus, when you create a fully-provisioned 10GB volume, 10GB of space is reserved for this volume in the cluster; if you also select **2-Way Replication**, 20 GB of space (10 GB x 2) would be reserved. The **Full Provisioning** option ensures that the full space requirement is reserved for a volume within the storage cluster.
- **Thin Provisioning**
  With **Thin Provisioning**, you reserve less space in the storage cluster than that presented to XenServer hosts. Thus, when a thinly-provisioned 10GB volume is created, only 1GB of space is initially reserved for this volume; however, a 10GB volume is presented to the host. If you were also to select **2-Way Replication**, 2GB of space (1 GB x 2) would initially be reserved for this volume.

As the initial 1GB reservation becomes almost consumed by writes, additional space is reserved from available space on the storage cluster. As more and more writes occur, the full 10GB of space will eventually be reserved.

Benefits of thin provisioning

The key advantage of using thin provisioning is that it minimizes the initial storage footprint during deployment. As your needs change, you can increase the size of the storage cluster by adding storage nodes to increase the amount of space available, creating a cost-effective, pay-as-you-grow architecture.

When undertaking a project to consolidate servers through virtualization, you typically find under-utilized resources on the bare-metal server; however, storage tends to be over-allocated. Now, XenServer's resource virtualization approach means that storage can also be consolidated in clusters; moreover, thin provisioning can be selected to optimize storage utilization.

As your storage needs grow, you can add storage nodes to increase performance and capacity – a single, simple GUI operation is all that is required to add a new node to a management group and storage cluster. HP SAN/iQ storage software automatically redistributes your data based on the new cluster size, immediately providing additional space to support the growth of thinly-provisioned volumes. There is no need to change VM configurations or disrupt access to live data volumes.

However, there is a risk associated with the use of thin provisioning. Since less space is reserved on the SAN than that presented to XenServer hosts, writes to a thinly-provisioned volume may fail if the SAN should run out of space. To minimize this risk, SAN/iQ software monitors utilization and issues warnings when a cluster is nearly full, allowing you to plan your data growth needs in conjunction with thin provisioning. Thus, to support planned storage growth, it is a best practice to configure e-mail alerts, Simple Network Management Protocol (SNMP) triggers, or CMC storage monitoring so that you can initiate an effective response prior to a full-cluster event. Should a full-cluster event occur, writes requiring additional space cannot be accepted and will fail until such space is made available, effectively forcing the SR offline.

In order to increase available space in a storage cluster, you have the following options:

- Add another storage node to the SAN
- Delete other volumes
- Reduce the volume replication level

Note
Reducing the replication level or omitting replication frees up space; however, the affected volumes would become more prone to failure.

- Replicate volumes to another cluster and then delete the originals

Note
After moving volumes to another cluster, you would have to reconfigure XenServer host access to match the new SR.

Adding a storage node to a cluster may be the least disruptive option for increasing space without impacting data availability.

This section has provided guidelines and best practices for configuring a new iSCSI volume. The following section describes how to configure a XenServer host.

Configuring a XenServer Host

This section provides guidelines and best practices for configuring a XenServer host so that it can communicate with an HP StorageWorks P4000 SAN, ensuring that the storage bandwidth for each VM is optimized. For example, since XenServer iSCSI SRs depend on the underlying network configuration, you can maximize availability by bonding network interfaces; in addition, you can create a dedicated storage network to achieve predictable storage throughput for VMs. After you

have configured a single host in a resource pool, you can scale up with additional hosts to enhance VM availability.

The sample SRs configured below utilize the iSCSI volumes described in the previous section.

Guidelines are provided for the following tasks:

- Synchronizing time between XenServer hosts
- Setting up networks and configuring network bonding
- Connecting to iSCSI volumes in the SR iSCSI Storage Repositories that will be created utilizing the HP StorageWorks iSCSI volumes created in the previous section
- Creating a VM on the SR and best practices implemented ensuring that each virtual machine maximizes its available iSCSI storage bandwidth

The section ends with a summary.

Synchronizing time

A server's BIOS provides a local mechanism for accurately recording time; in the case of a XenServer host, its VMs also use this time.

By default, XenServer hosts are configured to use local time for time stamping operations. Alternatively, a network time protocol (NTP) server can be used to manage time for a management group rather than relying on local settings.

Since XenServer hosts, VMs, applications, and storage nodes all utilize event logging, it is considered a best practice – particularly when there are multiple hosts – to synchronize time for the entire virtualized environment via an NTP server. Having a common time-line for all event and error logs can aid in troubleshooting, administration, and performance management.

Note
Configurations depend on local resources and networking policy.
NTP synchronization updates occur every five minutes.
If you do not set the time zone for the management group, Greenwich Mean Time (GMT) is used.

NTP for HP StorageWorks P4000 SAN
NTP Server configuration can be found on a Management Group's Time tab. Common time for all event and error logs can aid in troubleshooting, administration, and performance management. A preferred NTP Server of 0.pool.ntp.org is used in this example shown in Figure 5.

Figure 5. Turning on NTP using the CMC

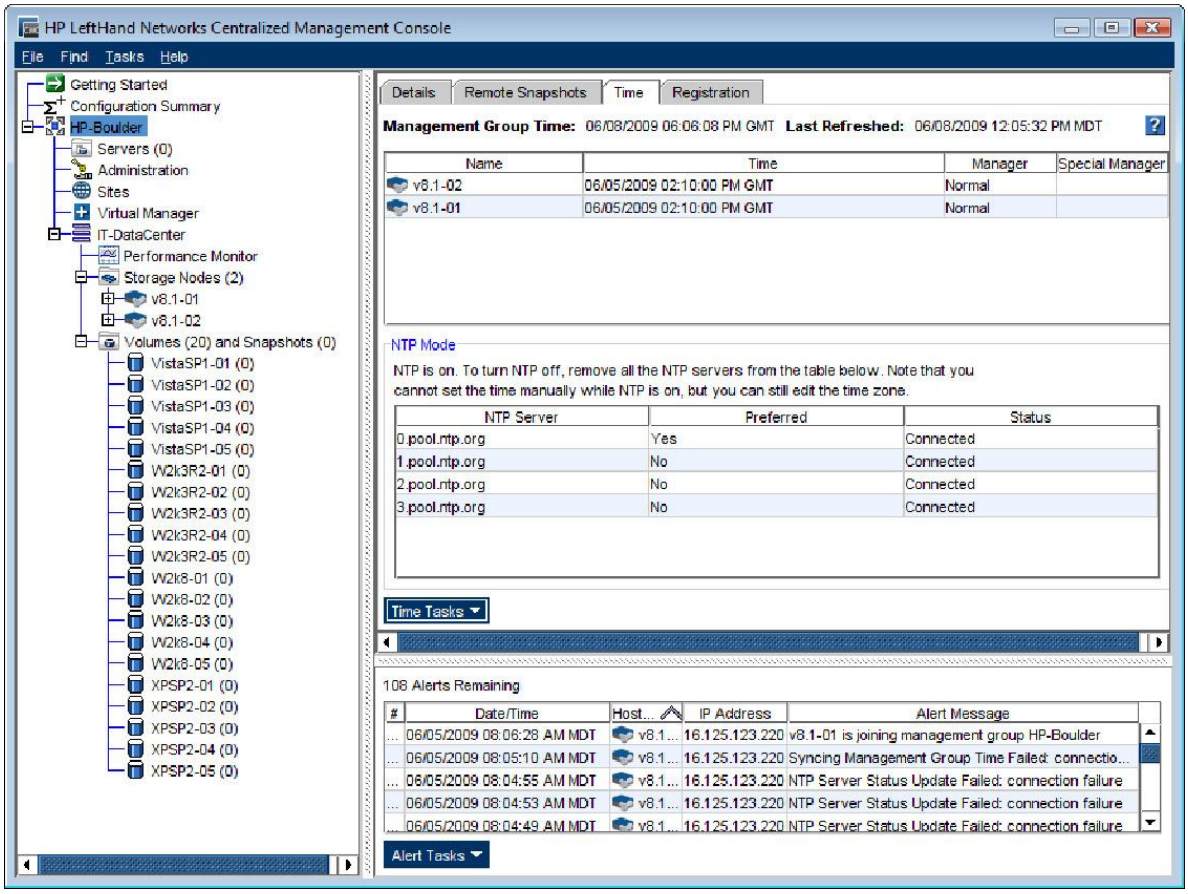

NTP for XenServer

Although NTP Server configuration may be performed during a XenServer installation, the console may also be used post installation. Within XenCenter, highlight the XenServer and select the Console tab. Enable NTP using xsconsole. Enable NTP as shown in Figure 6.

Figure 6. Turning on NTP using the XenServer xsconsole

Network configuration and bonding

Network traffic to XenServer hosts may consist of the following types:

- XenServer management
- VM LAN traffic
- iSCSI SAN traffic

Although a single physical network adapter can accommodate all these traffic types, its bandwidth would have to be shared by each. However, since it is critically important for iSCSI SRs to perform predictably when serving VM storage, it is considered a best practice to dedicate a network adapter to the iSCSI SAN. Furthermore, to maximize the availability of SAN access, you can bond multiple network adapters to act as a single interface, which not only provides redundant paths but also increases the bandwidth for SRs. If desired, you can create an additional bond for LAN connectivity.

Note
XenServer supports source-level balancing (SLB) bonding.

It is a best practice to ensure that the network adapters configured in a bond have matching physical network interfaces so that the appropriate failover path can be configured. In addition, to avoid a SPOF at a common switch, multiple switches should be configured for each failover path to provide an additional level of redundancy in the physical switch fabric.

You can create bonds using either XenCenter or the XenServer console, which allows you to specify more options and must be used to set certain bonded network parameters for the iSCSI SAN. For example, the console must be used to set the **disallow-unplug** parameter to **true**.

Example

In the following example, six separate network links are available to a XenServer host. Of these, two are bonded for VM LAN traffic and two for iSCSI SAN traffic. In general, the procedure is as follows:

1. Ensure there are no VMs running on the particular XenServer host.
2. Select the host in XenCenter and open the **Network** tab, as shown in Figure 7.

A best practice for the networks is to add a meaningful description to each network in the description field.

Figure 7. Select host in XenCenter and open Network tab.

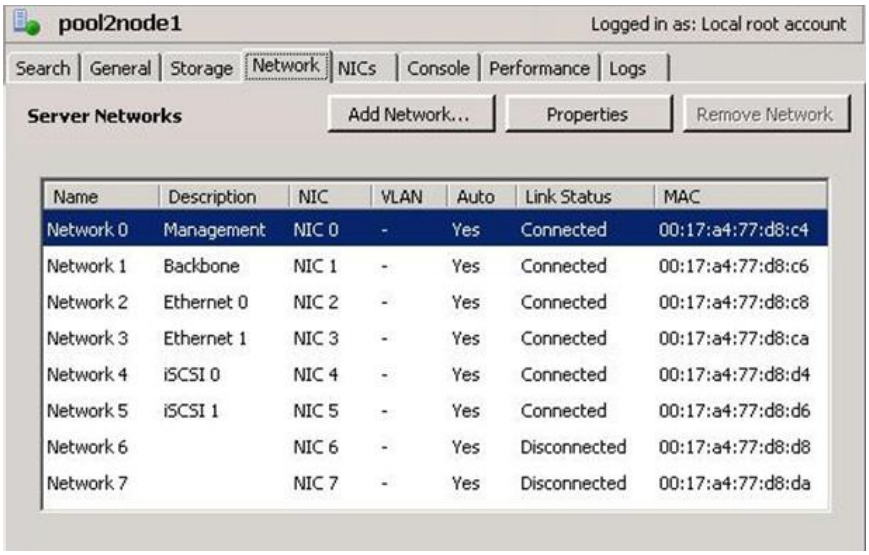

3. Select the **NICs** tab and click the **Create Bond** button. Add the interfaces you wish to bond, as shown in Figure 8.

Figure 8. Bonding network adapters **NIC 4** and **NIC 5**

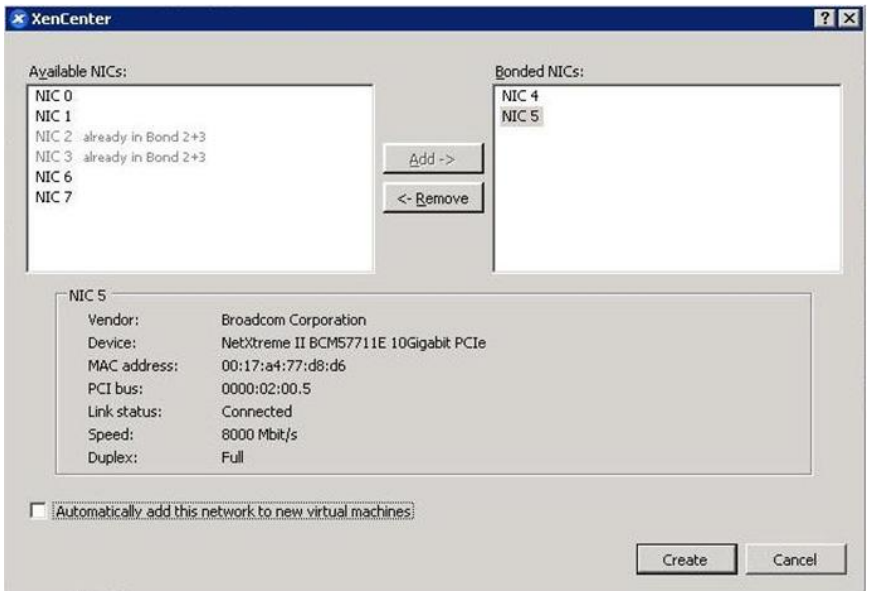

Figure 8 shows the creation of a network bond consisting of NIC 4 and NIC 5 to connect the host to the iSCSI SAN and, thus, the SRs that are common to all hosts. NIC 2 and NIC 3 had already been bonded to form a single logical network link for Ethernet traffic.

The network in this example consists of a class C subnet of **255.255.255.0** with a network address of **1.1.1.0**. No gateway is configured. IP addressing is set using the **pif-reconfigure-ip** command.

4. As shown in Figure 9, select **Properties** for each bonded network; rename **Bond 2+3** to **Bond 0** and rename **Bond 4+5** to **Bond 1**; and enter appropriate descriptions for these networks.

Figure 9. Renaming network bonds

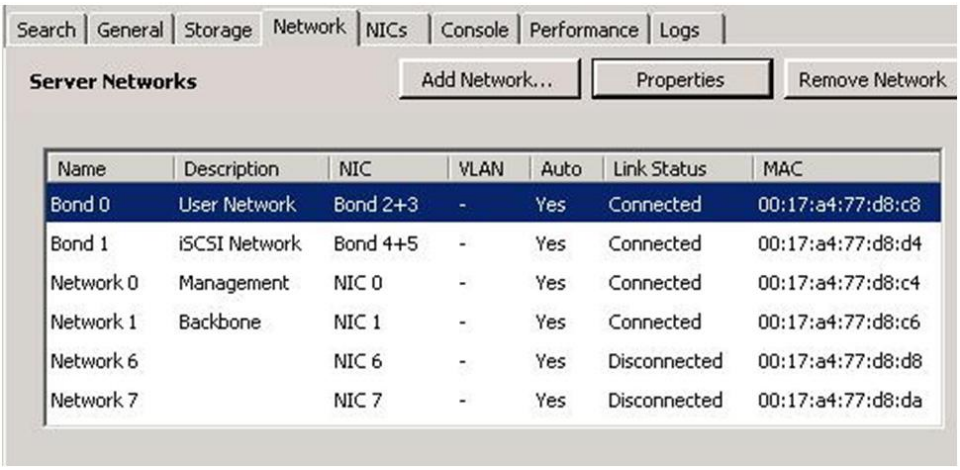

The iSCSI SAN Bond 1 interface is now ready to be used. In order for the bond's IP address to be recognized, you can reboot the XenServer host; alternatively, use the **host-management-reconfigure** command.

Connecting to an iSCSI volume

While HP StorageWorks iSCSI volumes were created in a previous section, no access was assigned to those volumes.

Before a volume can be recognized by a XenServer host as an SR, you must use the CMC to define the authentication method to be assigned to this volume. The following authentication methods are supported on XenServer hosts:

- **IQN** – You can assign a volume based on an IQN definition. Think of this as a one-to-one relationship, with one rule for one host.
- **CHAP** – Challenge Handshake Authentication Protocol (CHAP) provides a mechanism for defining a user name and secret password credential to ensure that access to a particular iSCSI volume is appropriate. Think of this as a one-to-many relationship, with one rule for many hosts.

XenServer hosts only support one -way CHAP credential access.

Determining or changing the host's IQN

Each XenServer host has been assigned a default IQN that you can change, if desired; in addition, each volume has been assigned a unique iSCSI name during its creation. Although specific naming is not required within an iSCSI SAN, all IQNs must be unique from initiator to target.

A XenServer host's IQN can be found via XenCenter's **General** tab, as shown in Figure 10. Here, the IQN of XenServer host **XenServer-55b-02** is **iqn.2009-06.com.example:e834bedd**.

Figure 10. Determining the IQN of a particular XenServer host

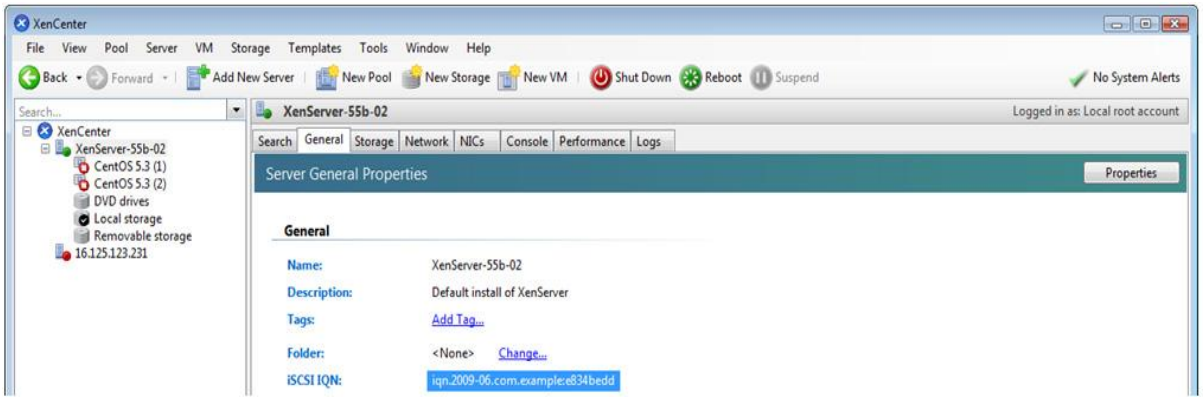

If desired, you can use the **General** tab's **Properties** button to change the host's IQN, as shown in Figure 11.

Figure 11. Changing the host's IQN

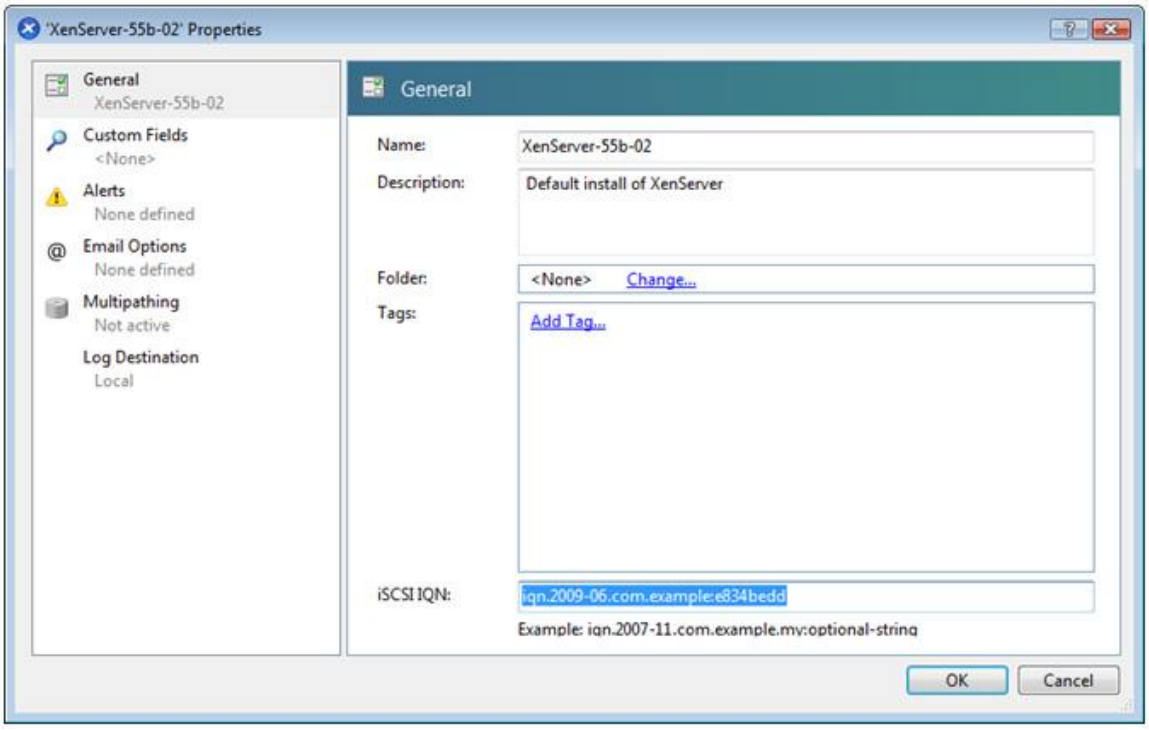

Note
Once you have used the CMC to define an authentication method for an iSCSI volume, if the host's IQN changes, you must update accordingly.

Alternatively, you can update a host's IQN via CMC's command-line interface (CLI). Use the **host-param-set command**.

Note
The host's Universally Unique Identifier (UUID) must be specified.

Verify the change using the **host-param-list** command.

Specifying IQN authentication

This subsection talks about an "authentication method" that is given a name.

Before specifying the authentication method for a volume, you can use the CMC to determine its IQN. In this example, SR access is created for volume **XPSP2-01,** which has an IQN of **iqn.2003-10.com.lefthandnetworks:hp-boulder:55:xpsp2-01** (obtained by highlighting this volume in the **Details** tab, as shown in Figure 12.)

Figure 12. Obtaining the IQN of volume **XPSP2-01**

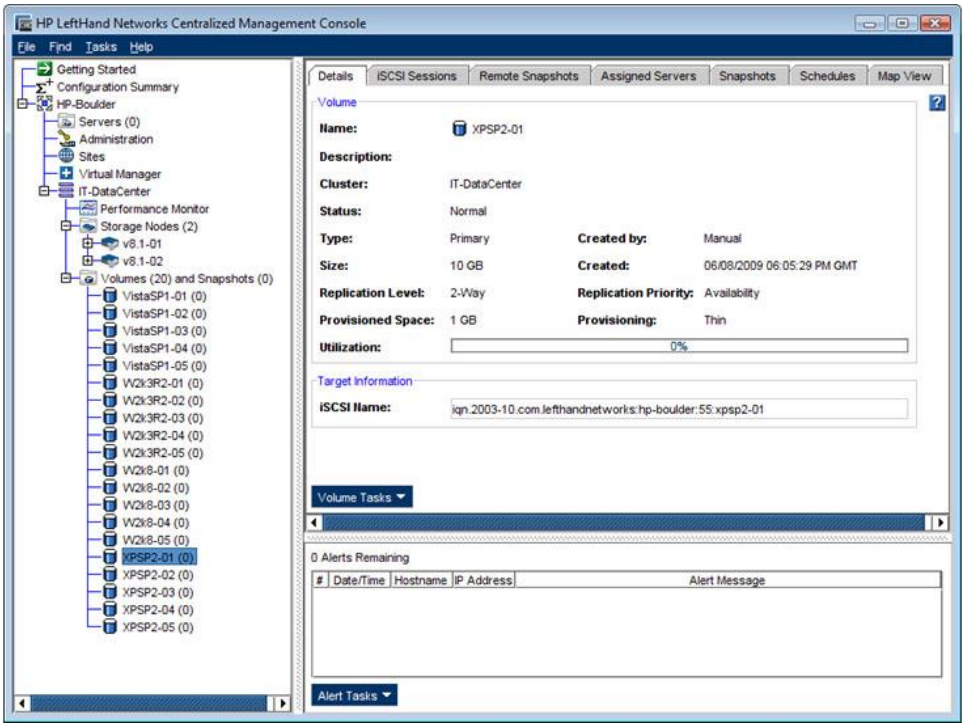

Use the following procedure:

1. Under **HP-Boulder**, highlight the **Servers (0)** selection. Note that the currently defined authentication rule method is currently zero (0).
2. To obtain the **New Server** dialog box (as shown in Figure 13), either right-click on **Servers (0)** and select **New Server→Select Server Tasks→New Server** or utilize **Tasks→Server→New Server**.

Figure 13. New Server dialog box

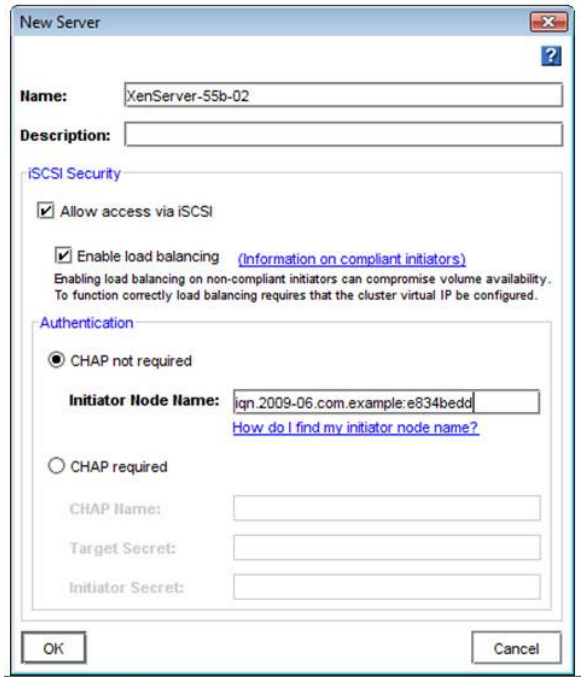

3. Enter the name **XenServer-55b-02**. Note that you can choose any name; however, matching the XenServer host name to the authentication method name implies the relationship between the two and makes it easier to assign iSCSI volumes in the CMC.

   Check **Allow access via iSCSI**.

   Check **Enable load balancing**.

   Under **CHAP not required**, enter the IQN of the host (**iqn.2009-06.com.example:e834bedd**) in the **Initiator Node Name** field.

4. After you have created the **XenServer-55b-02**, you can assign volumes and snapshots.

   Under the **Volumes and Snapshots** tab, select **Tasks→Assign and Unassign Volumes→Snapshots**. Alternatively, select **Tasks→Volume→Edit Volume→Selecting the Volume→Basic tab→Assign and Unassign Servers**. The former option focuses on assigning volumes and snapshots to a particular server (Figure 14); the latter on assigning servers to a particular volume (Figure 15).

   Assign access for volume **XPSP2-01** to the **XenServer-55b-02**.

   Specify access as **None**, **Read**, or **Read/Write**.

Figure 14. Assigning volumes and snapshots to server **XenServer-55b-02**

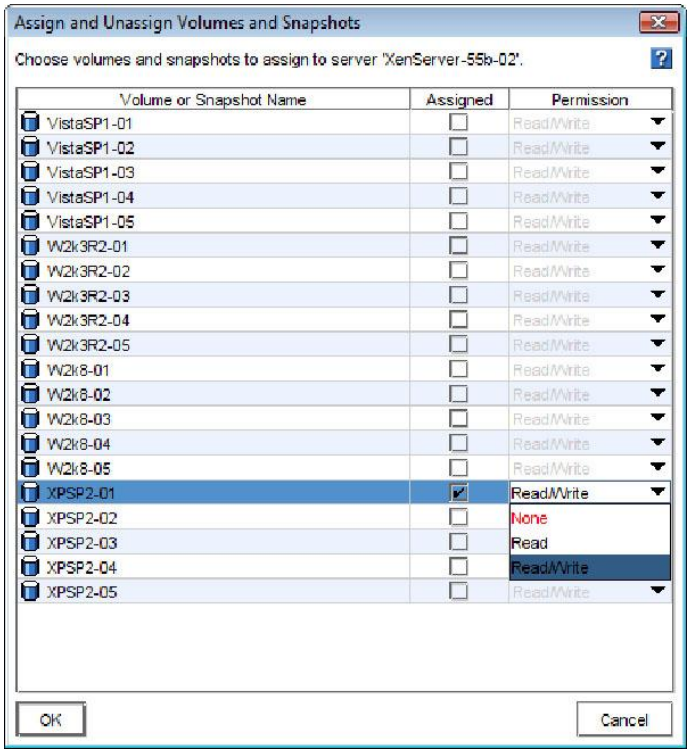

Figure 15. Assigning servers to volume **XenServer-55b-02**

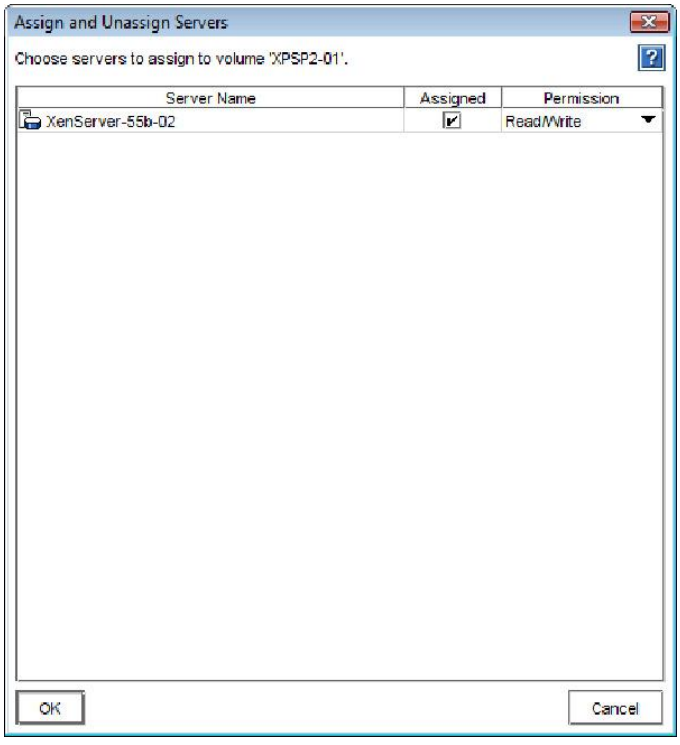

Creating an SR

Now that the XenServer host has been configured to access an iSCSI volume target, you can create a XenServer SR. You can configure an SR from HP StorageWorks SAN targets using LVM over iSCSI or LVM over HBA.

Note
LVM over HBA connectivity is beyond the scope of this white paper.

In this example, the IP address of host **XenServer-55b-02** is **1.1.1.230**; the virtual IP address of the HP StorageWorks iSCSI SAN cluster is **1.1.1.225**.

Use the following procedure to create a shared-LVM SR:

1. In XenCenter, select **Storage Repository** or, with XenCenter 5.5, **New Storage**.
2. Under **Virtual disk storage**, select **iSCSI** to create a shared-LVM SR, as shown in Figure 16. Select **Next**.

Figure 16. Selecting the option to create a shared-LVM SR

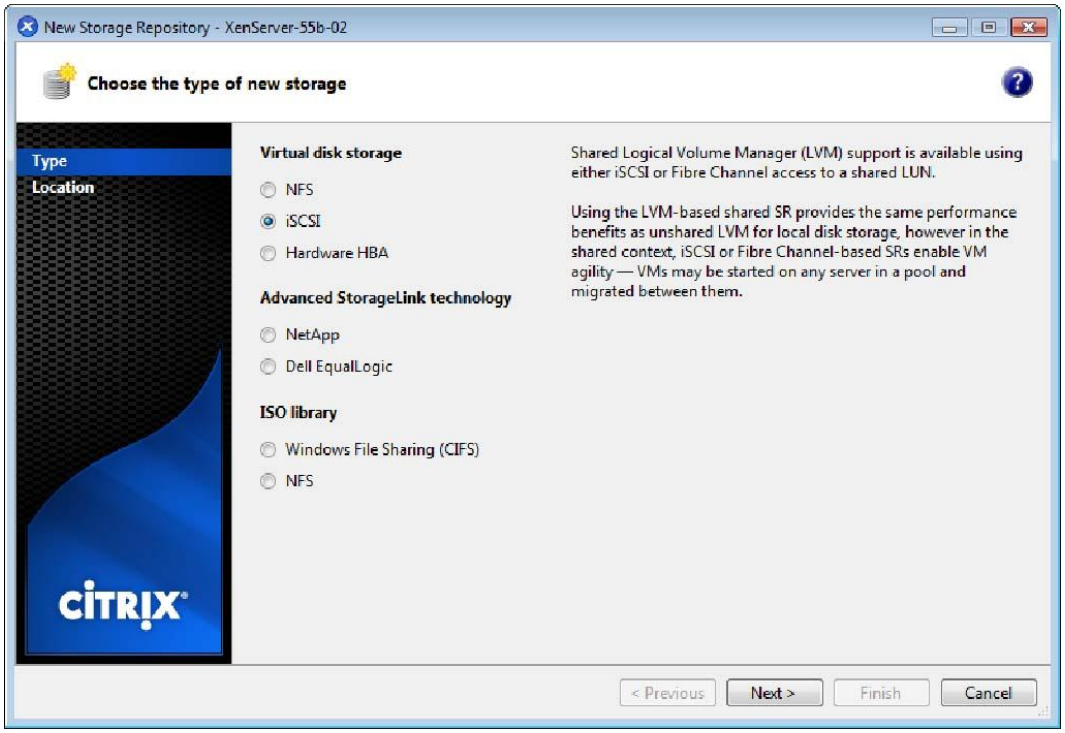
New Storage Repository - XenServer-55b-02
Choose the type of new storage
Type
Location
Virtual disk storage
NFS
iSCSI
Hardware HBA
Advanced StorageLink technology
NetApp
Dell EqualLogic
ISO library
Windows File Sharing (CIFS)
NFS
Shared Logical Volume Manager (LVM) support is available using either iSCSI or Fibre Channel access to a shared LUN.
Using the LVM-based shared SR provides the same performance benefits as unshared LVM for local disk storage, however in the shared context, iSCSI or Fibre Channel-based SRs enable VM agility — VMs may be started on any server in a pool and migrated between them.
CITRIX
< Previous
Next >
Finish
Cancel

3. Specify the name of and path to the SR. For clarity, the name **XPSP2-01** is used to match the name of the iSCSI volume and of the VM that will be created later.

Figure 17. Naming the SR **XPSP2-01**

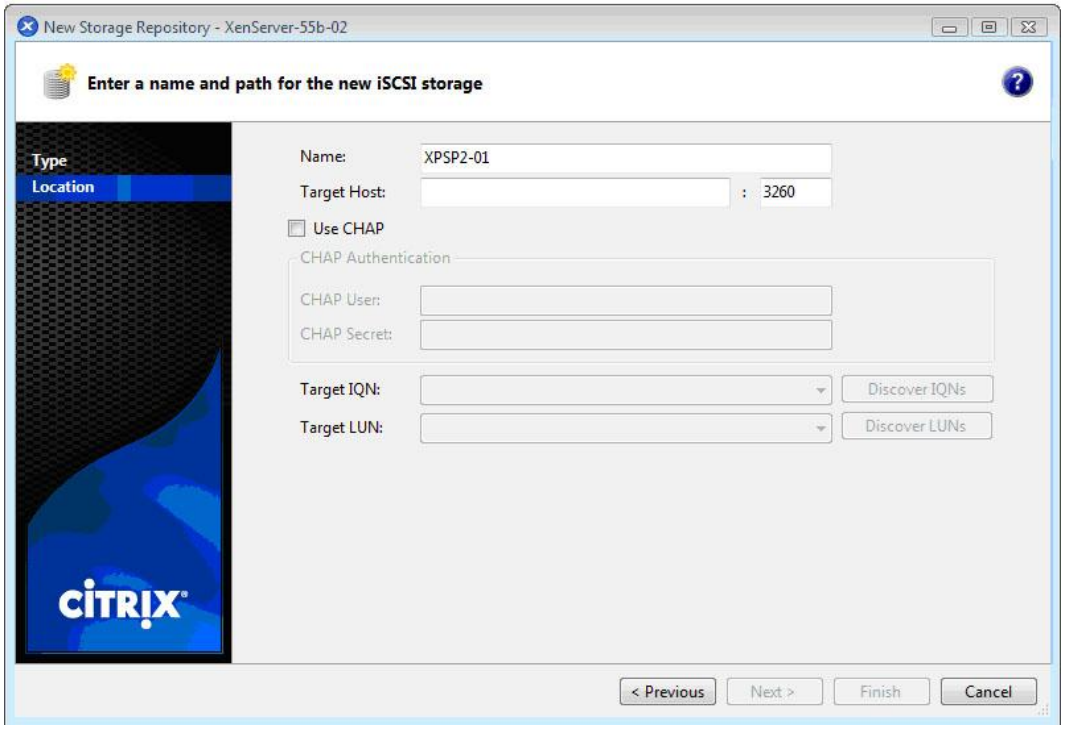

4. As shown in Figure 18, specify the target host for the SR as **1.1.1.225** (the virtual IP address of the HP StorageWorks iSCSI SAN cluster).

   Next, select **Discover IQNs** to list visible iSCSI storage targets in a drop-down list. Match the **Target IQN** value to the IQN of volume **XPSP2-01** as shown in the CMC.

   Select **Discover LUNs;** then specify **LUN 0** as the **Target LUN**, forcing iSCSI to be presented at LUN 0 for each unique target IQN.

   Select **Finish** to complete the creation of the new SR.

Figure 18. Specify the target IQN and LUN

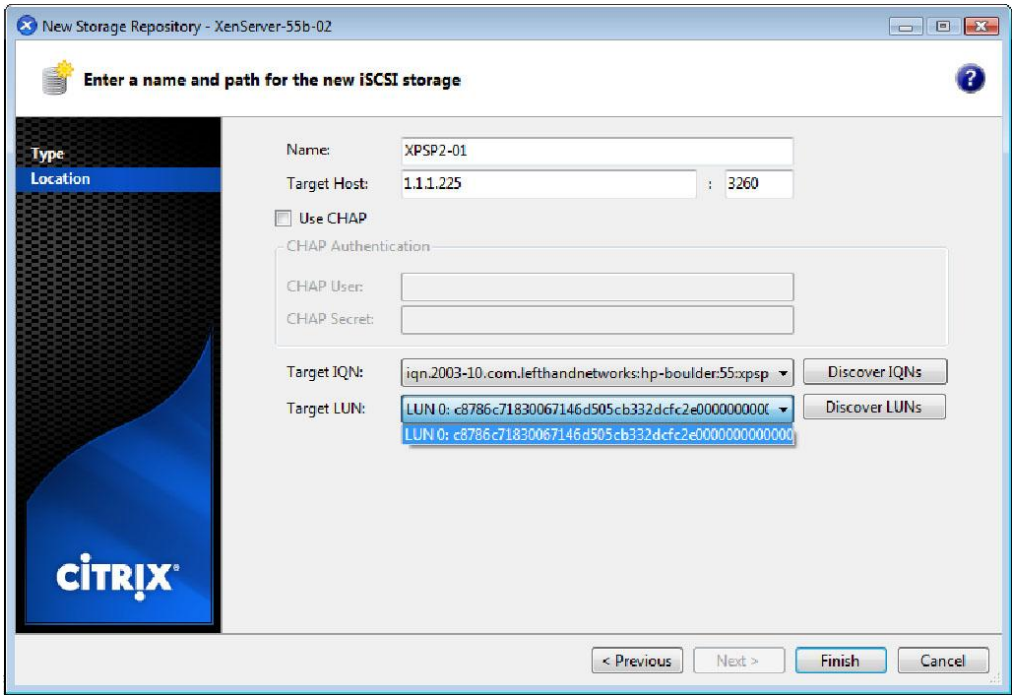

5. For an LVM over iSCSI SR, raw volumes must be formatted before being presented to the XenServer host for use as VM storage.

   As shown in Figure 19, any data on a volume that is not in an LVM format will be lost during the format operation.

   After the format is complete, the SR will be available and enumerated in XenCenter under the XenServer host, as shown in Figure 20.

Figure 19. Warning that the format will destroy data on the volume

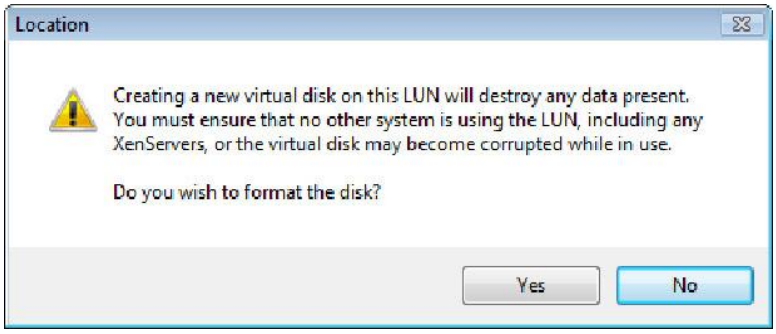
Location
Creating a new virtual disk on this LUN will destroy any data present. You must ensure that no other system is using the LUN, including any XenServers, or the virtual disk may become corrupted while in use.
Do you wish to format the disk?
Yes
No

Figure 20. Verifying that the enumerated SR is shown as available in XenCenter

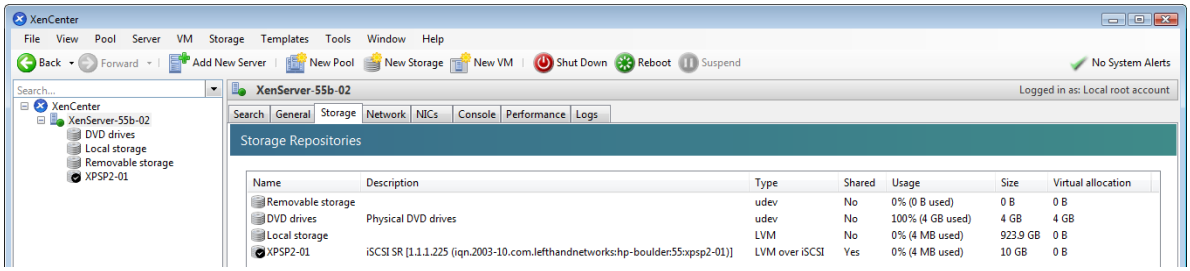
XenCenter
File View Pool Server VM Storage Templates Tools Window Help
Back Forward Add New Server New Pool New Storage New VM Shut Down Reboot Suspend
No System Alerts
XenCenter
XenServer-55b-02
DVD drives
Local storage
Removable storage
XPSP2-01
XenServer-55b-02
Logged in as: Local root account
Search General Storage Network NICs Console Performance Logs
Storage Repositories
Name Description Type Shared Usage Size Virtual allocation
Removable storage udev No 0% (0 B used) 0 B 0 B
DVD drives Physical DVD drives udev No 100% (4 GB used) 4 GB 4 GB
Local storage LVM No 0% (4 MB used) 923.9 GB 0 B
XPSP2-01 iSCSI SR [1.1.1.225 (iqn.2003-10.com.lefthandnetworks:hp-boulder:55:xpsp2-01)] LVM over iSCSI Yes 0% (4 MB used) 10 GB 0 B

Creating a VM on the new SR

Use the following procedure to create a VM on the SR you have just created.

1. From XenCenter's top menu, select **VM→New**.
2. Select the desired operating system template for the new VM. In this example, the VM will be running Microsoft® Windows® XP SP2.
3. For consistency, specify the VM's name as **XPSP2-01**, as shown in Figure 21.

Figure 21. Naming the new VM

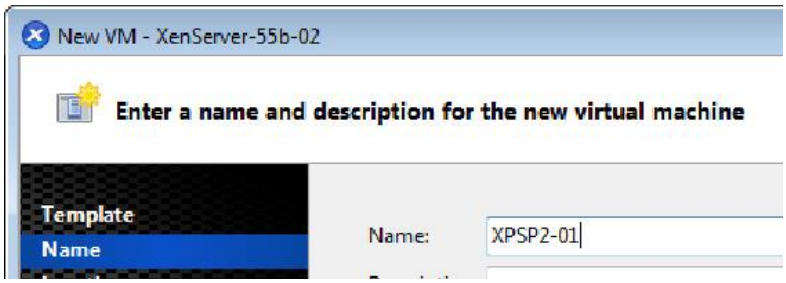

4. Select the type of installation media to be used, either **Physical DVD Drive** (used in this example) or **ISO Image**.

Note
A XenServer host can create an ISO SR library or import a Server Message Block (SMB)/Common Internet File System (CIFS) share. For more information, refer to your XenServer documentation.

5. Specify the number of virtual CPUs required and the initial memory allocation for the VM.

   These values depend on the intended use of the VM. For example, while the default memory allocation of 512MB is often sufficient, you may need to select a different value based on the particular VM's usage or application. If you do not allocate sufficient memory, paging to disk will cause performance contention and degrade overall XenServer performance.

   A typical Windows XP SP2 VM running Microsoft Office should perform adequately with 768MB.

   Thus, to optimize XenServer performance, it is a best practice to understand a VM's application and use case before its deployment in a live environment.

6. Increase the size of the virtual disk from 8GB (default) to 9GB, as shown in Figure 22. While the iSCSI volume is 10GB, some space is consumed by LVM SR overhead and is not available for VM use.

Note
The virtual disk presented to the VM is stored on the SR.

Figure 22. Changing the size of the virtual disk presented to the VM

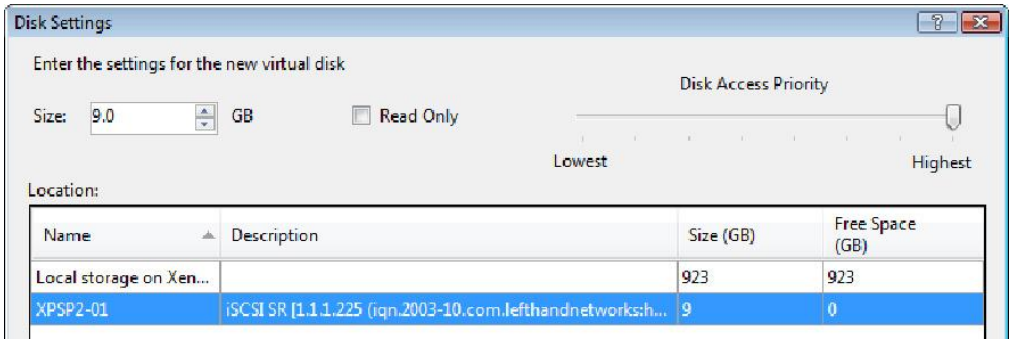

7. Allocate a single network interface – **interface0** – to the VM, which connects the VM to the **bond0** network for LAN access.
8. Do not select the option to start the VM automatically.
9. In order to install the operating system, insert Windows XP SP2 installation media into the XenServer host's local DVD drive. After a standard installation, the VM is started.

After Windows XP SP2 has been installed, XenCenter displays the started VM with an icon showing a green circle with a white arrow, as shown in Figure 23.

Note that the name of the VM, **XPSP2-01**, matches that of the SR associated with it, which is a best practice intended to provide clarity while configuring the environment.

The first SR is designated as the default and is depicted by an icon showing a black circle and a white check mark. Note that the default SR is used to store virtual disks, crash dump data, and images of suspended VMs.

Figure 23. Verifying that the new VM and SR are shown in XenCenter

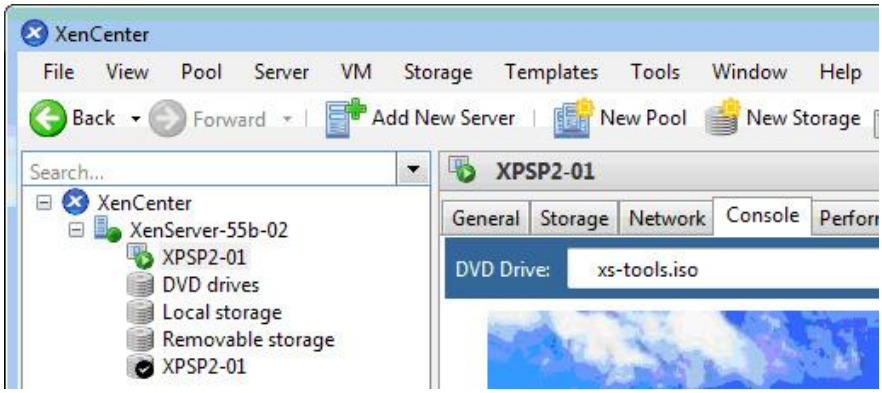

Summary

In the example described above, the following activities were performed:

- A XenServer host was configured with high-resiliency network bonds for a dedicated SAN and a LAN.
- An HP StorageWorks P4000 SAN was configured as a cluster of two storage nodes.
- A virtualized 10GB iSCSI volume, **XPSP2-01**, was configured with Network RAID and allocated to the host.
- A XenServer SR, **XPSP2-01**, was created on the iSCSI volume.
- A VM, **XPSP2-01**, with Windows XP SP2 installed, was created on a 9GB virtual disk on the SR.

Figure 24 outlines this configuration, which can be managed as follows:

- The XenCenter management console is installed on a Windows client that can access the LAN.
- The VM's local console is displayed with the running VM. Utilizing the resources of the XenServer host, the local console screen is transmitted to XenCenter for remote viewing.

Figure 24. The sample environment

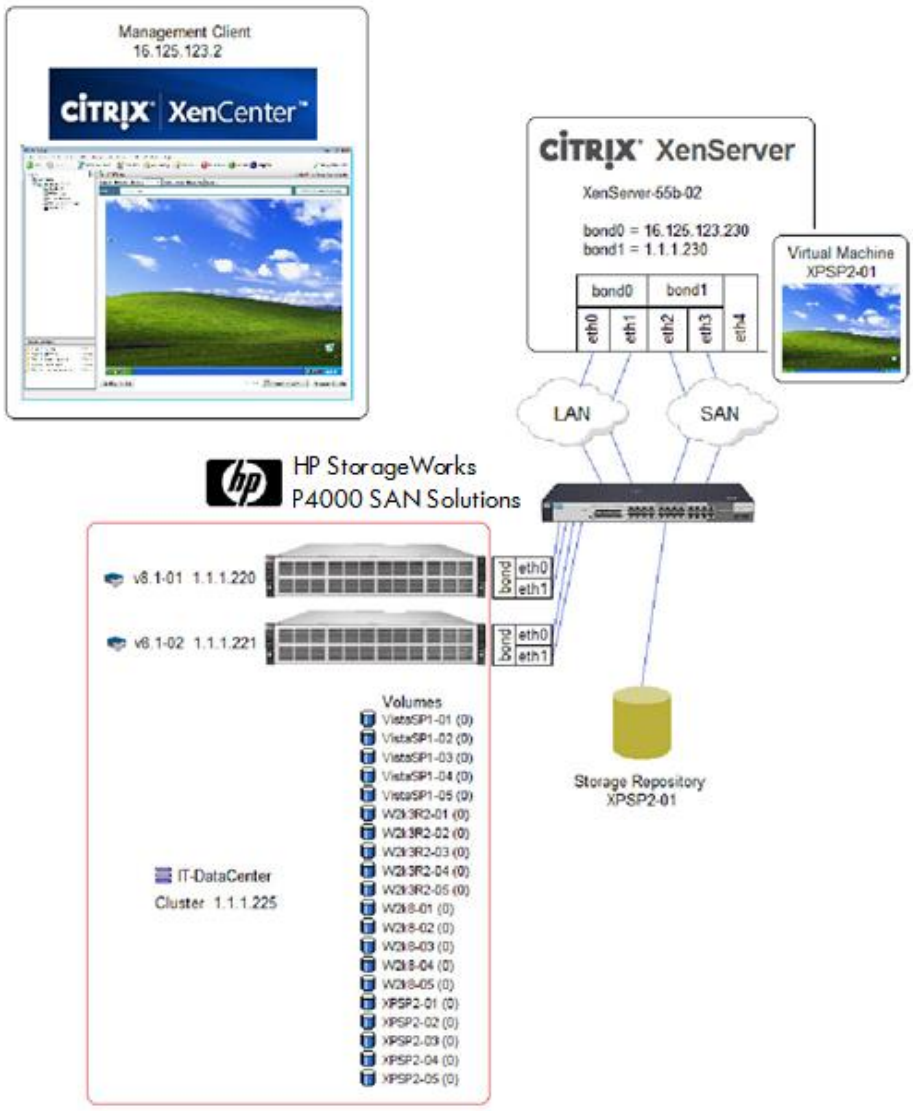

Configuring for high availability

After virtualizing physical servers that had been dedicated to particular applications and consolidating the resulting VMs on a XenServer host, you must ensure that the host will be able to run these VMs. Designing for high availability means that the components of this environment – from servers to storage to infrastructure – must be able to fail without impacting the delivery of associated applications or services.

For example, HP StorageWorks P4000 SANs offer the following features that can increase availability and avoid SPOFs:

- Hardware-based RAID for the storage node
- Multiple network interfaces bonded together

- Network RAID across the cluster of storage nodes.

XenServer host machines also deliver a range of high-availability features, including:

- Resource pools of XenServer hosts
- Multiple network interfaces bonded together
- The use of external, shared storage by VMs
- VMs configured for high availability

To help eliminate SPOFs from the infrastructure, network links configured as bonded pairs can be connected to separate physical switches.

In order to survive a site failure, XenServer resource pools can be configured as a stretch cluster across multiple sites, with HP StorageWorks P4000 SANs maintaining multi-site, synchronously-replicated SAN volumes that can be accessed at these same sites.

Enhancing infrastructure availability

The infrastructure that has been configured in earlier sections of this white paper (shown in Figure 24) has an SPOF – the network switch. By adding a second physical switch and connecting each bonded interface to each switch, as shown in Figure 25, this SPOF can be removed by adding a network switch.

Implementing switch redundancy for the LAN and SAN increases the resiliency of both VMs and applications.

Figure 25. Adding a network switch to remove a SPOF from the infrastructure

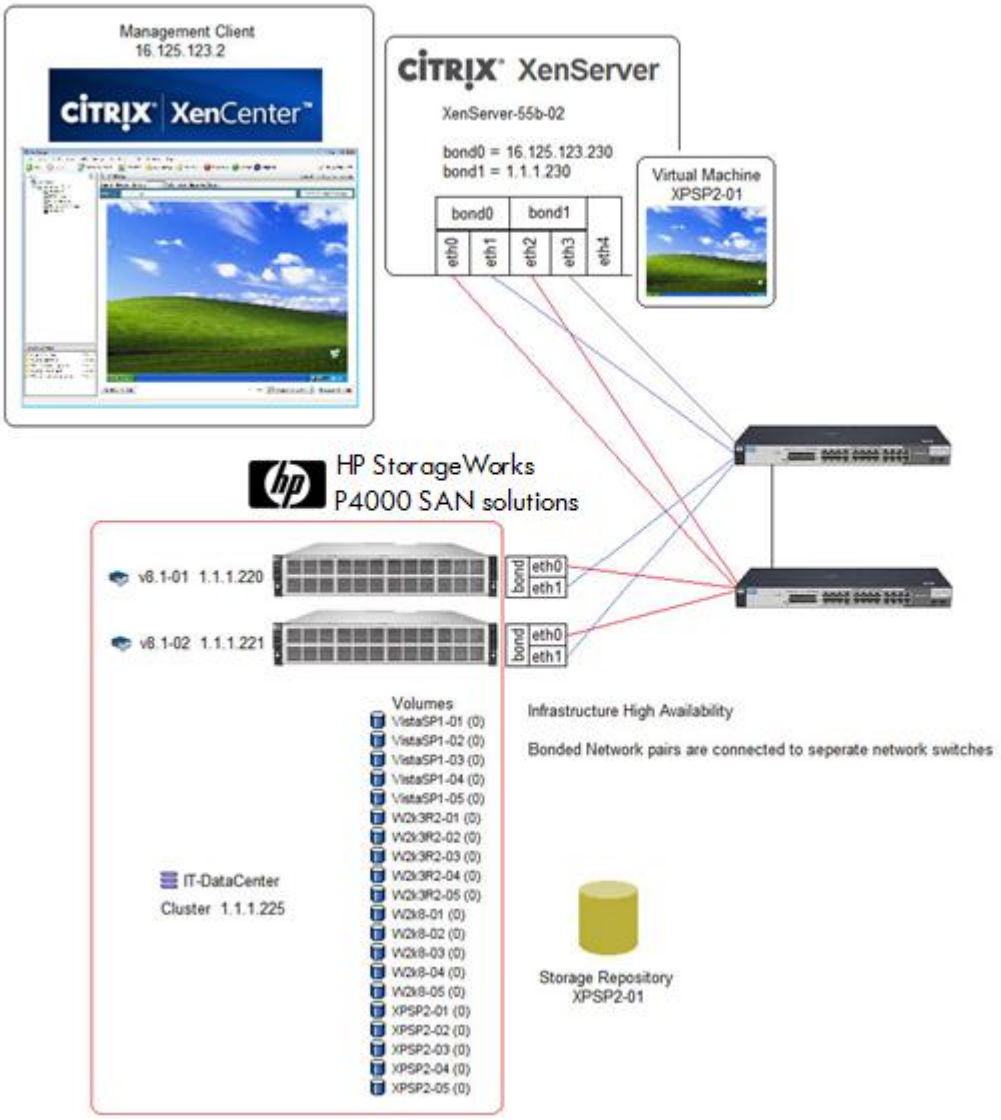

Note the changes to the physical connections to each switch – in order to be able to survive a switch failure in the infrastructure, each link in each bond must be connected to a separate switch.

Configuration

Consider the following when configuring your infrastructure:

- **HP** StorageWorks P4000 SAN **bonds** – You must configure the networking bonds for adaptive load balancing (ALB); Dynamic LACP (802.3ad) cannot be supported across multiple switch fabrics.
- **XenServer host bonds** – SLB bonds can be supported across multiple switches.

Implementing Network RAID for SRs

HP StorageWorks P4000 SANs can enhance the availability of XenServer SRs. The resiliency provided by clustering storage nodes, with each node featuring multiple controllers, multiple network interfaces, and multiple disk drives, can be enhanced by implementing Network RAID to protect the

logical volumes. With Network RAID, which is configurable on a per-volume basis, data blocks are written multiple times to multiple nodes. In the example shown in Figure 26, Network RAID has been configured with Replication Level 2, guaranteeing that a volume remains available despite the failure of multiple nodes.

Figure 26. The storage cluster is able to survive the failure of two nodes

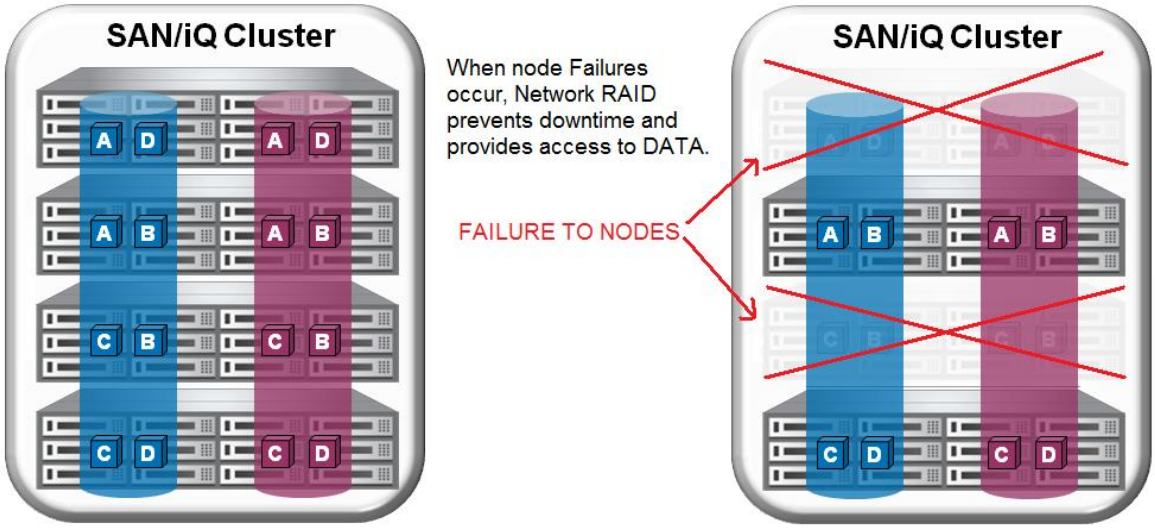
SAN/iQ Cluster
When node Failures occur, Network RAID prevents downtime and provides access to DATA.
FAILURE TO NODES
SAN/iQ Cluster

Configuring Network RAID

You can set the Network RAID configuration for a particular volume either during its creation or by editing the volume's properties on the **Advanced** tab, as shown in Figure 27. You can update Network RAID settings at any time without taking the volume down.

The following Network RAID options are provided:

- **2-Way** – Up to one adjacent node failures in a cluster as seen above
- **3-Way** – Up to two adjacent node failures in a cluster
- **4-Way** – Up to three adjacent node failures in a cluster

The number of nodes in a par**ticul**ar cluster determines which and how many nodes can fail, based on the selected Network RAID configuration.

Figure 27. Configuring Network RAID for a particular volume

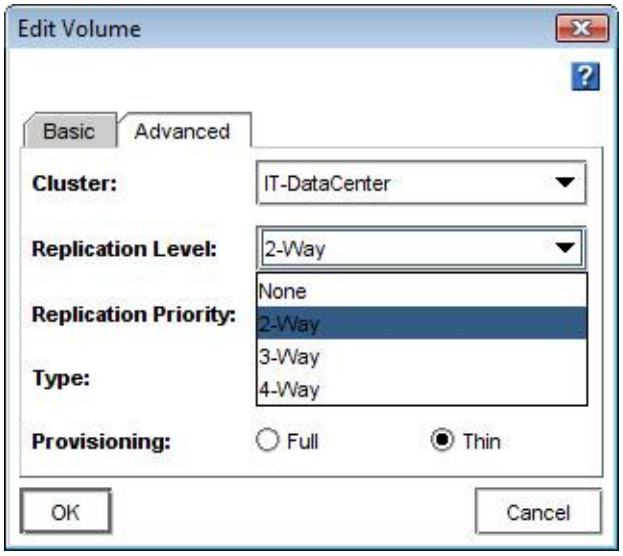

Pooling XenServer hosts

Multiple XenServer hosts can be deployed to support VMs, with each host utilizing its own resources and acting as an individual virtualization platform. To enhance availability, however, consider creating a XenServer host resource pool (that is, a group of similarly-configured XenServer hosts working together as a single entity with shared resources, as shown in Figure 28). A pool of up to 16 XenServer hosts can be used to run and manage a large number of VMs.

Figure 28. A XenServer host resource pool with two host machines

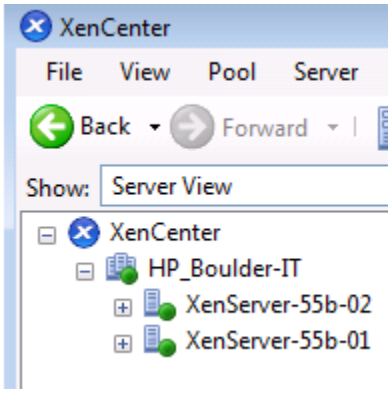

Key to the success of a host resource pool is the deployment of SAN-based, shared storage, providing each host with equal access that appears to be local.

With shared storage, VMs can be configured for high availability. In the event of a XenServer host failure, a VM would leverage Citrix XenMotion functionality to automatically migrate from the failed host to another host in the pool.

From XenCenter, you can discover multiple XenServer hosts that are similarly configured with resources.

Configuring VMs for high availability

You can use XenServer's High Availability (HA) feature to enhance the availability of a XenServer resource pool. When this option is enabled, XenServer continuously monitors the health of all hosts in a resource pool; in the event of a host failure, specified VMs would automatically be moved to a healthy host.

In order to detect the failure of a host, XenServer uses multiple heartbeat-detection mechanisms through shared storage as well as network interfaces.

Note
It is a best practice to create bonded interfaces in HA configurations to maximize host resiliency, thereby preventing false positives of component failures.

You can enable and configure HA using the Configure HA wizard.

Creating a heartbeat volume

Since XenServer HA requires a heartbeat mechanism within the SR, you should create a special HP StorageWorks iSCSI volume for this purpose. This heartbeat volume must be accessible to all members of the resource pool and must have a minimum size of 356MB.

It is a best practice to name the heartbeat volume after the resource pool, adding "HeartBeat" for clarity. Thus, in this example, resource pool **HP_Boulder-IT** includes a 356MB iSCSI volume named **HP-Boulder-IT-HeartBeat** that can be accessed by both hosts, **XenServer-55b-01** and **XenServer-55b-02**, as shown in Figure 29.

Figure 29. A volume named **HP-Boulder-IT-HeartBeat** has been added to the resource pool

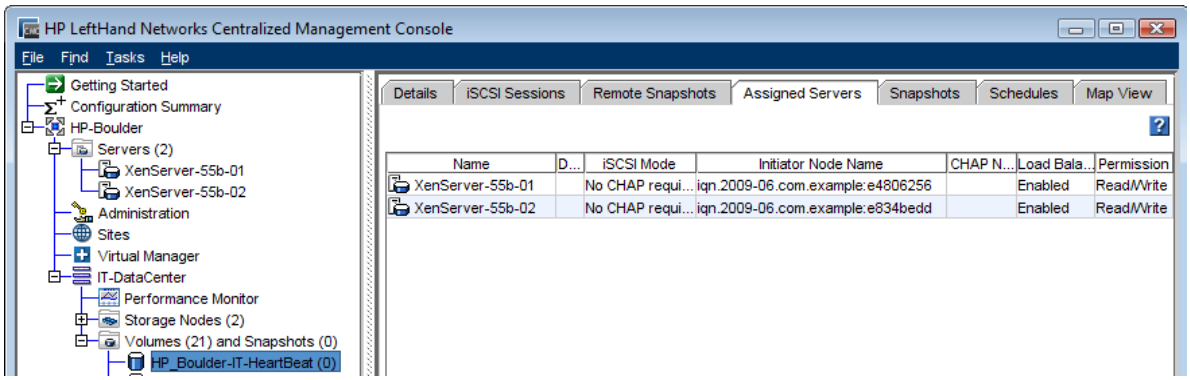

You can now use XenCenter to create a new SR for the heartbeat volume. For consistency, name the SR **HP_Boulder-IT-HeartBeat**.

As shown in Figure 30, the volume appears in XenCenter with 356MB of available space; 4MB is used for the heartbeat and 256MB for pool master metadata.

Figure 30. The properties of **HP-Boulder-IT-HeartBeat**

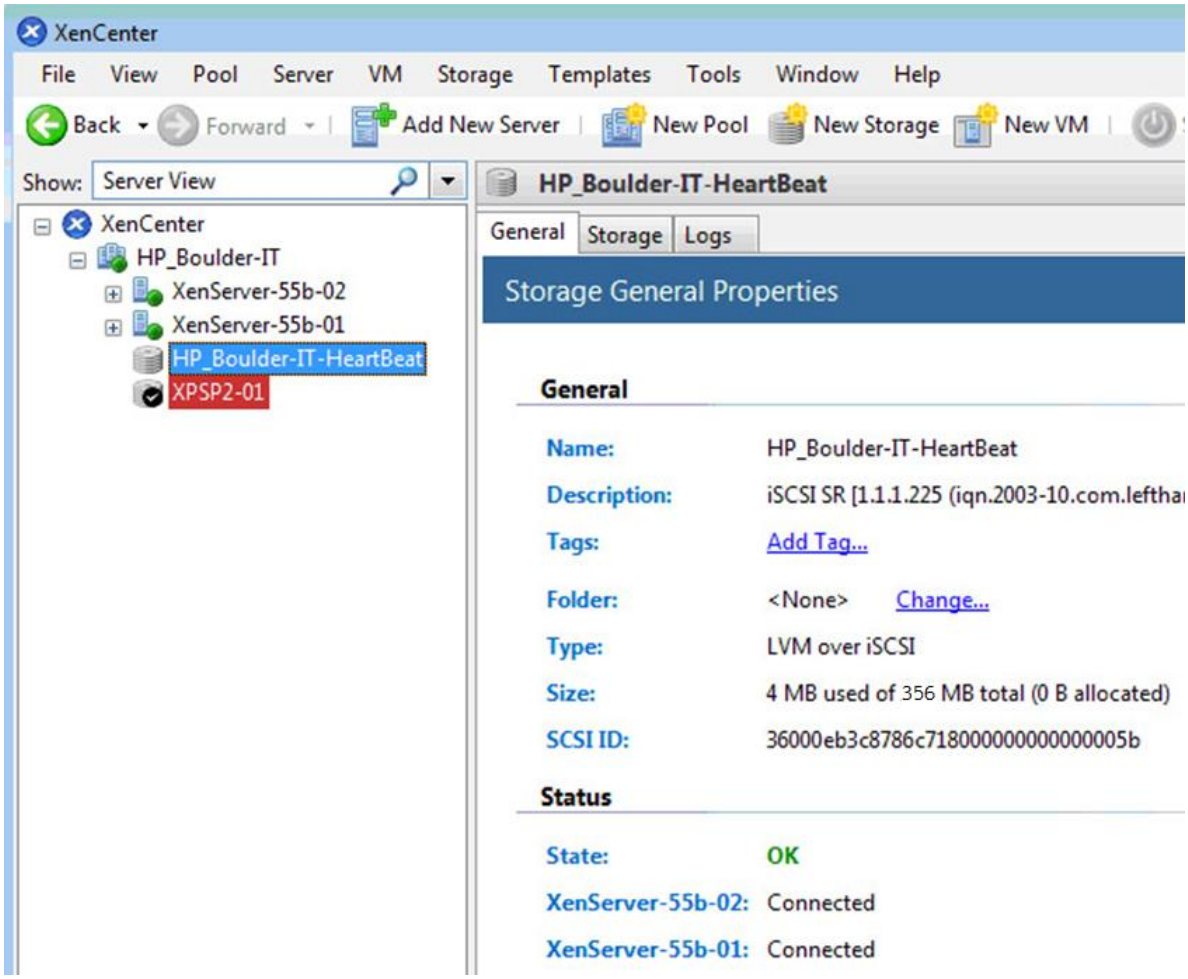

Configuring the resource pool for HA

XenServer HA maintains a failover plan that defines the response to the failure of one or more XenServer hosts. To configure HA functionality for a resource pool, select the **Pool** menu option in XenCenter and click **High Availability**. The **Configure HA** wizard now guides you through the setup, allowing you, for example, to specify the heartbeat SR, **HP_Boulder-IT-HeartBeat**, as shown in Figure 31.

As part of the setup, you can specify HA protection levels for VMs currently defined in the pool. Options include:

- **Protect** – The "Protect" setting ensures that HA has been enabled for VMs. Protected VMs are guaranteed to be restarted first if sufficient resources are available within the pool.
- **Restart if possible** – VMs with a "Restart if possible" setting are not guaranteed to be restarted following a host failure. However, if sufficient resources are available after protected VMs have been restarted, "Restart if possible" VMs will be restarted on a surviving pool member.
- **Do not restart** – VMs with a "Do not restart" setting are not restarted following a host failure.

If the resource pool does not have sufficient resources available, even protected VMs are not guaranteed to be restarted. XenServer will make additional attempts to restart VMs when the state of

the resource pool changes. For example, if you shut down non-essential VMs or add hosts to the pool, XenServer would make a fresh attempt to restart VMs. You should be aware of the following caveats:

- XenServer does not automatically stop or migrate running VMs in order to free up resources so that VMs from a failed host can be restarted elsewhere.
- If you wish to shut down a protected VM to free up resources, you must first disable its HA protection. Unless HA is disabled, shutting down a protected VM would trigger a restart.

You can also specify the number of server failures to be tolerated.

Figure 31. Configuring resource pool **HP-Boulder-IT** for HA

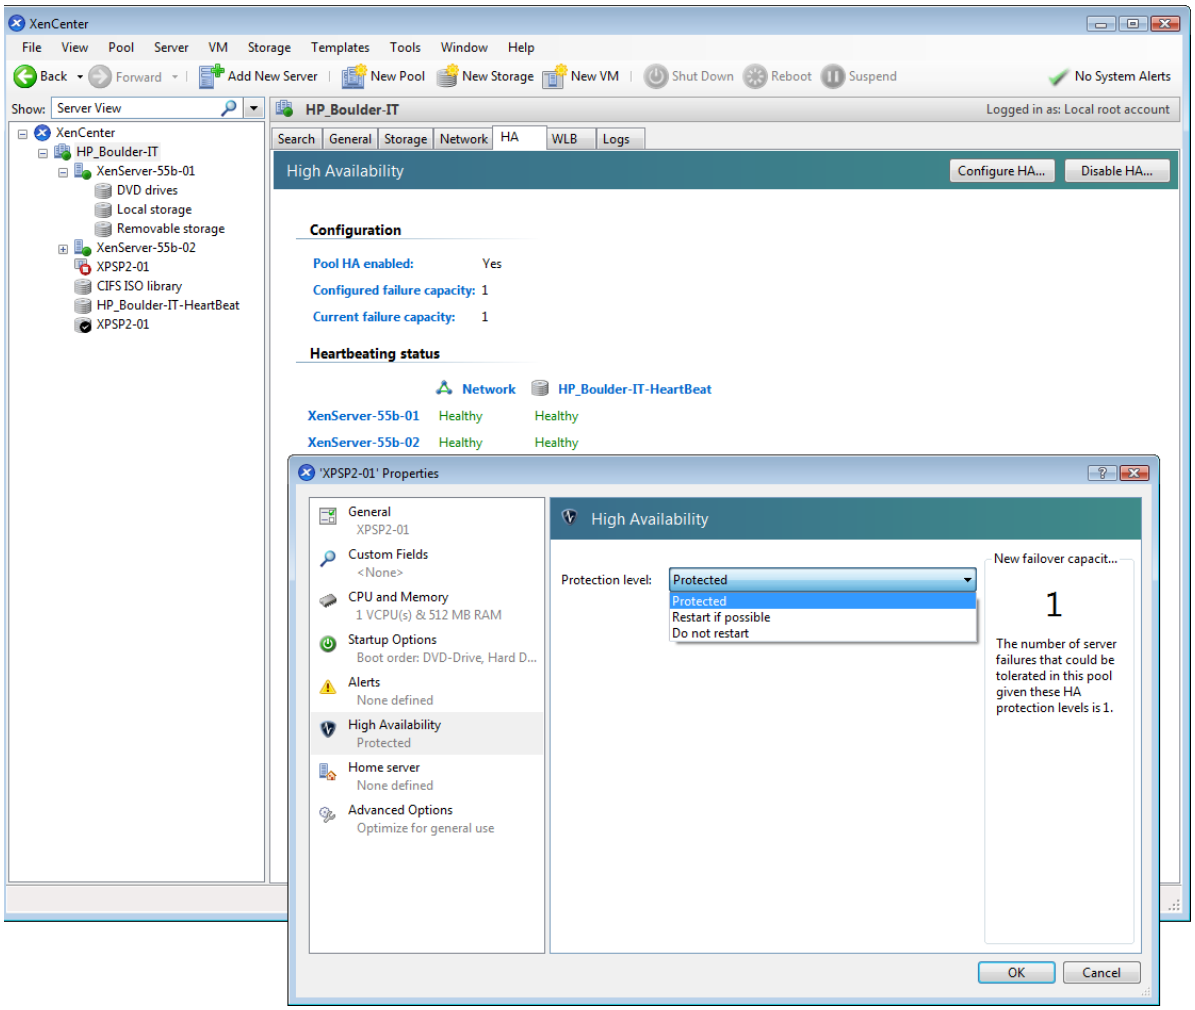

XenCenter provides a configuration event summary under the resource pool's **Logs** tab.

Following the HA configuration, you can individually tailor the settings for each VM using its **Properties** tab (select **Properties→High Availability**).

If you create additional VMs, the Configure HA can be used as a summary page for status for all high availability VMs.

Configuring multi-site high availability with a single cluster

If your organization deploys multiple data centers in close proximity, communicating over low-latency, high-bandwidth connections[4], you can stretch a resource pool between both sites. In this scenario, an entire data center is no longer a SPOF.

The stretched resource pool continuously constantly transfers pool status and management information over the network. Status information is also maintained on the 356MB iSCSI shared volume.

To give VMs the agility to start on any host on any site, SRs must be synchronously replicated and made available at each site. HP StorageWorks P4000 SANs support this scenario by allowing you to create a divided, replicated SAN cluster (known as a multi-site SAN), as shown in Figure 32.

Note
Since storage is being simultaneously replicated to both sites, you should ensure that appropriate bandwidth is available on the SAN.

Figure 32. Sample stretch pool configuration with multi-site SAN

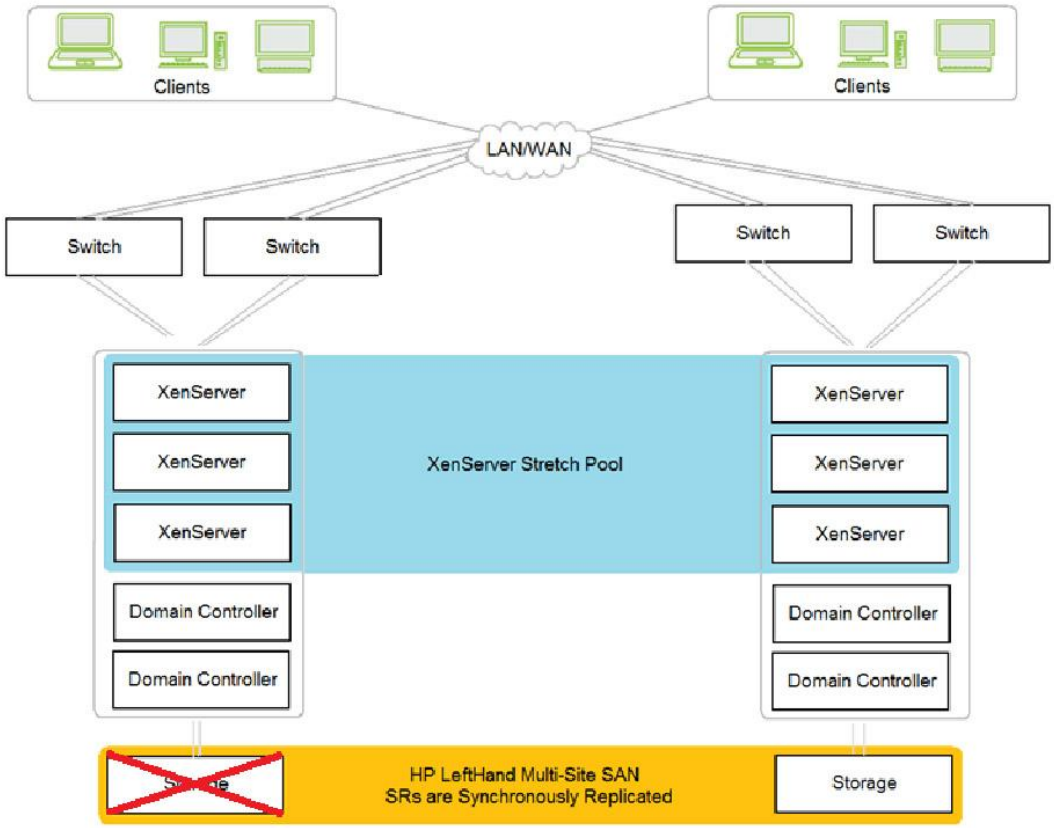

Requirements for a multi-site SAN include the following:

- Storage nodes must be configured in separate physical subnets at each site.
- Virtual IPs must be configured for the target portal for each subnet.

[4] The connection between sites needs to exhibit network performance similar to that of a single-site configuration.

- Appropriate physical and virtual networks exist at both sites.

Alternatively, the multi-site SAN feature can be implemented by correct physical node placement of single site cluster.

In a two-site implementation, you need an even number of storage nodes whether you have chosen a single-site cluster or multi-site SAN. Each site must contain an equal number of storage nodes.

In the case of a single-site SAN that is physically separated, it is critical that either the order of nodes in the cluster is appropriately created or the correct node is chosen from the order for physical separation. This requirement is straightforward in a two-node cluster; however, with a four-node cluster, you would deploy nodes 1 and 3 at Site A, while deploying nodes 2 and 4 at Site B. With the four-node cluster, volumes must, at a minimum, be configured for Replication Level 2; thus, even if Site A or B were to go down, storage nodes at the surviving site can still access the volumes required to support local VMs as well as protected VMs from the failed site.

You can find the order of storage nodes in the CMC by editing the cluster, as shown in Figure 33. If desired, you can change the order by highlighting a particular storage node and moving it up or down the list. Note that changing the order causes all volumes in the cluster to be restriped, which may take some time.

Figure 33. Editing the cluster

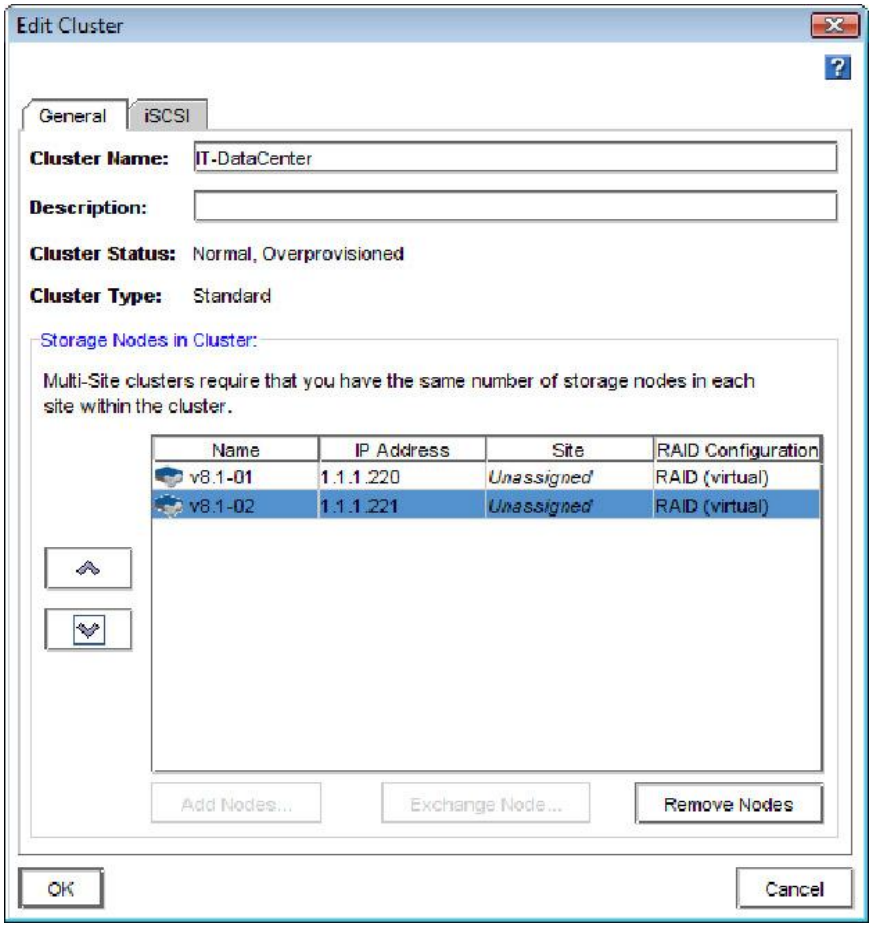

Note
It is a best practice to physically separate the appropriate nodes or ensure the order is valid before creating volumes.

Configuring multi-site high availability with multiple clusters

If multiple data centers are located at some distance from each other or the connections between them are high-latency, low–bandwidth, you should not stretch a XenServer resource pool between these sites. Since the resource pools would constantly need to transfer pool status and management information across the network, a stretch pool would be impractical in this case. Instead, you should create separate resource pools at each site, as shown in Figure 34.

Figure 34. Editing the cluster

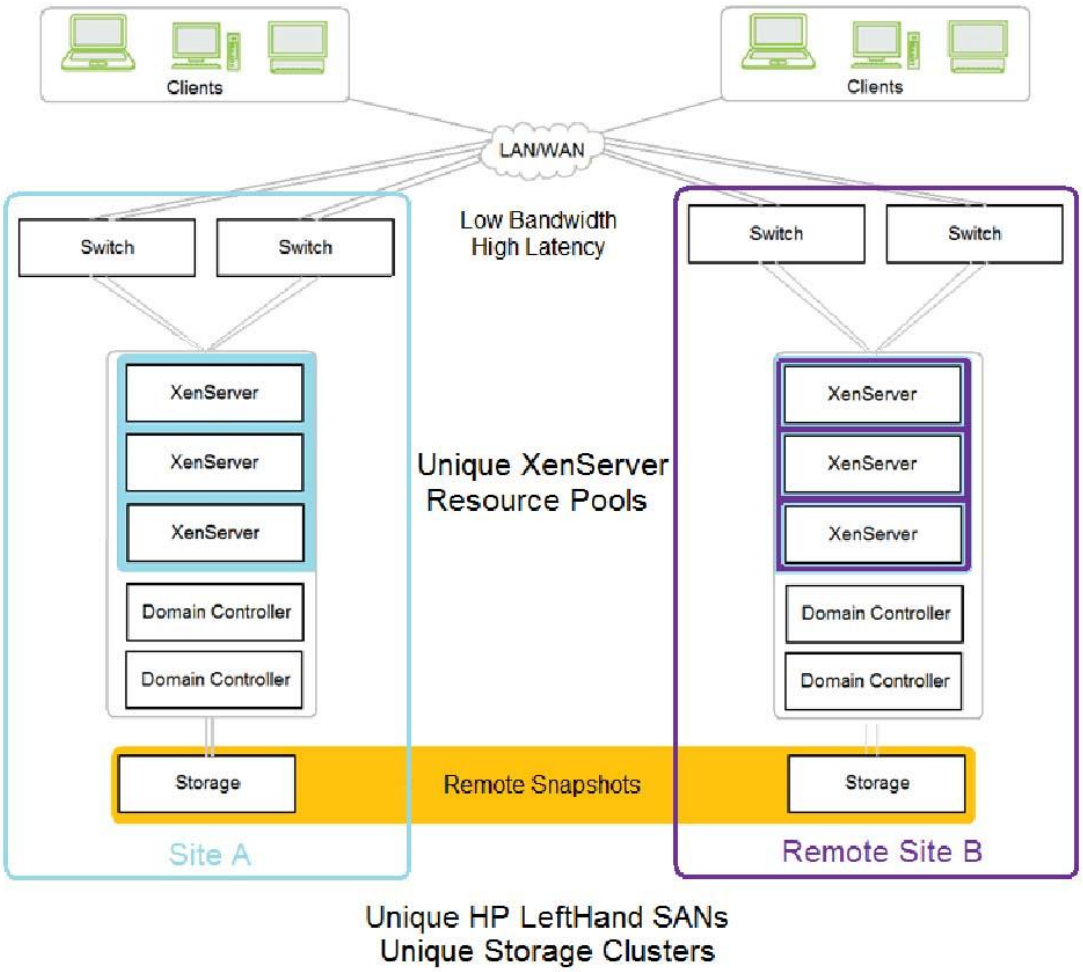

In the implementation shown in Figure 34, the remote site would be utilized in the event of the complete failure of the primary site (Site A). Resource pools at the remote site would be available to service mission-critical VMs from the primary site, delivering a similar level of functionality[5].

[5] You can expect some data loss due to the asynchronous nature of data snapshots.

When using an HP StorageWorks P4000 SAN, you would configure a management group at Site A. This management group consists of a cluster of storage nodes and volumes that serve Site A's XenServer resource pool; all VMs rely on virtual disks stored on SRs; in turn, the SRs are stored on highly-available iSCSI volumes. In order to survive the failure of this site, you must establish a remote snapshot schedule (as shown in Figure 35) to replicate these volumes to the remote site.

Figure 35. Creating a new remote snapshot

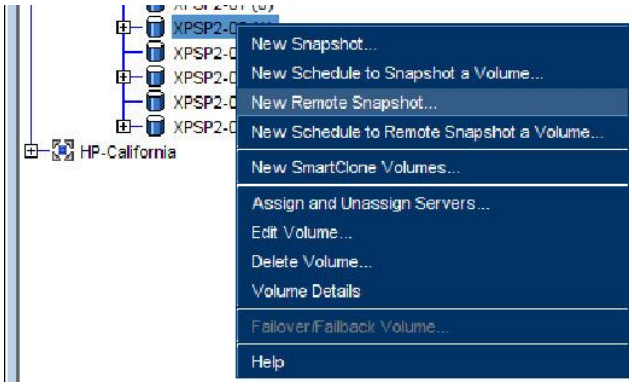

The initial remote snapshot is used to copy an entire volume to the remote site; subsequent scheduled snapshots only copy changes to the volume, thereby optimizing utilization of the available bandwidth. You can schedule remote snapshots based on the following criteria:

- Rate at which data changes
- Amount of bandwidth available
- Tolerability for data loss following a site failure

Remote snapshots can be performed sub-hourly, or less often (daily – weekly). These asynchronous snapshots provide a mechanism for recovering VMs at a remote site.

In any HA environment, you must make a business decision to determine which services to bring back online following a failover. Ideally, no data would be lost; however, even with sub-hourly (asynchronous) snapshots, some data from Site A may be lost. Since there are bandwidth limitations, choices must be made.

Creating a snapshot

Perform the following procedure to create a snapshot:

1. From the CMC, select the iSCSI volume you wish to replicate to the remote site.
2. Right-click on the volume and select **New Schedule to Remote Snapshot a Volume**.
3. Select the **Edit** button associated with **Start At** to specify when the schedule will commence, as shown in Figure 36.

Figure 36. Creating a new remote snapshot

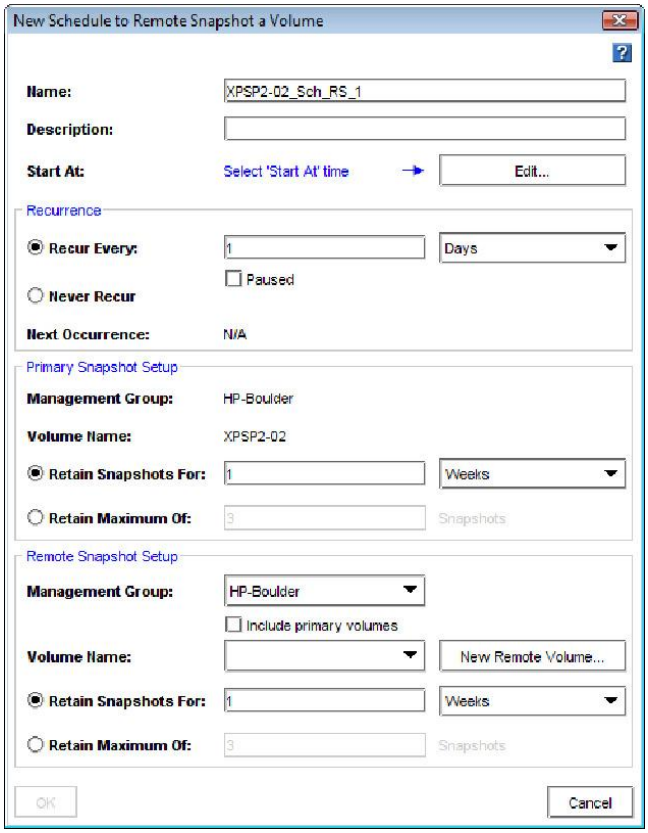

4. Set the **Recurrence** time (in minutes, hours, days, or weeks).
   Consider the following:
   – Ensure you leave enough time for the previous snapshot to complete.
   – Ensure there is adequate storage space at both sites.
   – Set a retention policy at the primary site based on a timeframe or snapshot count.
5. Select the **Management Group** for the remote snapshot.
6. Create a remote volume as the destination for the snapshot.
   Based on the convention used in this document, name the target **Remote-XPSP2-02**.
7. Set the replication level of **Remote-XPSP2-02**
8. Set the retention policy for the remote site. .

Depending on the scheduled start time, replication may now commence.

The first remote snapshot copies all your data – perhaps many terabytes – to the remote cluster.

To speed up the process, you can carry out this initial push on nodes at the local site, then physically ship these nodes to the remote site[6]. For more information, refer to the support documentation

After the initial push, subsequent remote snapshots are smaller – only transferring changes made since the last snapshot – and can make efficient use of a high-latency, low-bandwidth connection.

[6] The physical transfer of data – in this case, a storage cluster or SAN that may be carrying many terabytes of data – is known as **sneakernetting**.

Throttling bandwidth

Management groups support bandwidth throttling for data transfers, allowing you to manually configure bandwidth service levels for shared links.

In the CMC, right-click the management group, and select **Edit Management Group**. As shown in Figure 37, you can adjust bandwidth priority from **Fractional T1** (256 Kb/sec) to **Gigabit Ethernet** values.

Figure 37. Setting the remote bandwidth

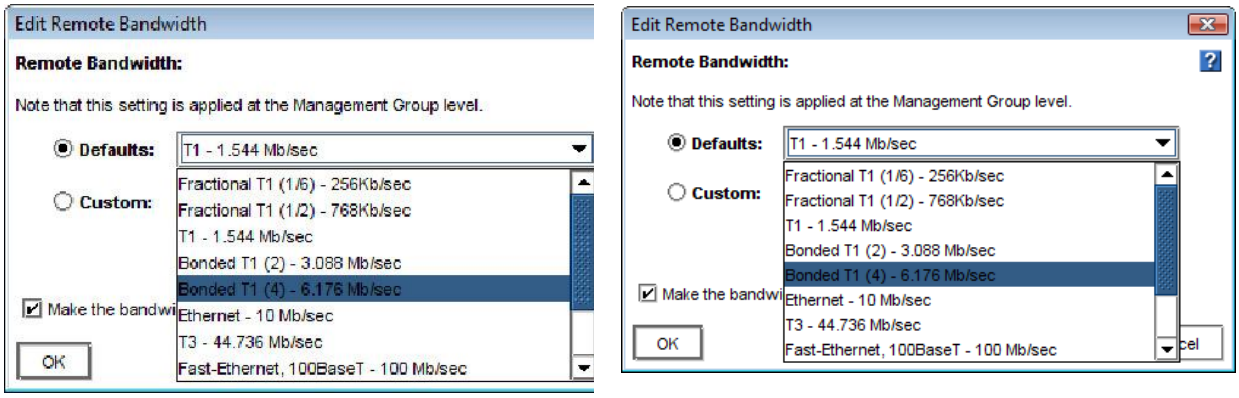
Edit Remote Bandwidth
Remote Bandwidth:
Note that this setting is applied at the Management Group level.
Defaults:
Custom:
T1 - 1.544 Mb/sec
Fractional T1 (1/6) - 256Kb/sec
Fractional T1 (1/2) - 768Kb/sec
T1 - 1.544 Mb/sec
Bonded T1 (2) - 3.088 Mb/sec
Bonded T1 (4) - 6.176 Mb/sec
Ethernet - 10 Mb/sec
T3 - 44.736 Mb/sec
Fast-Ethernet, 100BaseT - 100 Mb/sec
Make the bandwi
OK
Edit Remote Bandwidth
Remote Bandwidth:
Note that this setting is applied at the Management Group level.
Defaults:
Custom:
T1 - 1.544 Mb/sec
Fractional T1 (1/6) - 256Kb/sec
Fractional T1 (1/2) - 768Kb/sec
T1 - 1.544 Mb/sec
Bonded T1 (2) - 3.088 Mb/sec
Bonded T1 (4) - 6.176 Mb/sec
Ethernet - 10 Mb/sec
T3 - 44.736 Mb/sec
Fast-Ethernet, 100BaseT - 100 Mb/sec
Make the bandwi
OK

Starting VMs at the remote site

After a volume has been remotely copied, it contains the last completed snapshot, which is the most current data on the schedule.

In the event of a failure at Site A, volumes containing VM data at the remote site can be made primaries and the SRs reattached. VMs can then be started with the remote volume snapshots.

For more information on this process, refer to *Changing the Storage Repository and Virtual Disk UUID.*

Note
If Site A is down while the remote site is running the VMs, there is no need to change UUIDs.

Changing the direction of snapshots

As shown in Figure 38, after Site A has been restored, change the direction of remote snapshots to resynchronize snapshot data; you can then restart VMs at Site A. However, in the time taken to complete the snapshot and restart VMs, changes to the original data may have occurred; thus, data cannot be restored in this manner.

Note
With asynchronous snapshots performed in the other direction, it is assumed that a restoration may not include data that changed in the period between the last snapshot and the failure event.

You must accept the potential for data loss or use alternate methods for data synchronization.

Figure 38. Changing the direction

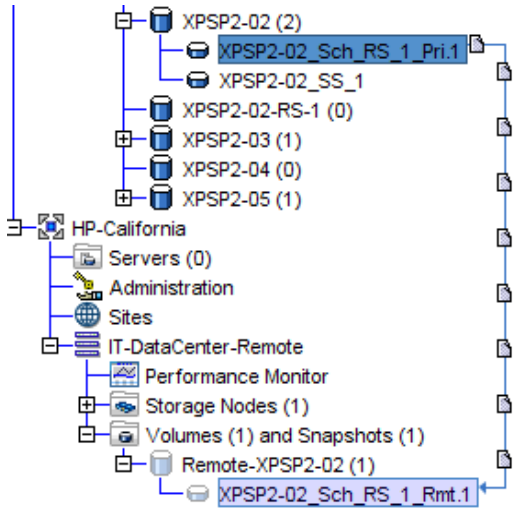

The CMC may be used with the Volume Failover/Failback Wizard.

Refer to the HP StorageWorks P4000 SAN User's Guide for additional information on documented procedures.

Disaster recoverability

Approaches to maximizing business continuity should rightly focus on preventing the loss of data and services. However, no matter how well you plan for disaster avoidance, you must also plan for disaster recovery.

Disaster recoverability encompasses the abilities to protect and recover data, and includes moving your virtual environment onto replacement hardware.

Since data corruption can occur in a virtualized environment just as easily as in a physical environment, you must predetermine restoration points that are tolerable to your business goals, along with the data you need to protect. You must also specify the maximum time it can take to perform a restoration, which is, effectively, downtime; it may be critical for your business to minimize this restoration time.

This section outlines different approaches to disaster recoverability. Although backup applications can be used within VMs, the solutions described here focus on the use of XenCenter tools and HP StorageWorks P4000 SAN features to back up data to disk and maximize storage efficiency. More information is provided on the following topics:

- Backing up configurations
- Backing up metadata
- Creating VM snapshots
- Copying a VM
- Creating SAN-based snapshots
- Rolling back a SAN-based snapshot

- Reattaching SRs

Backing up configurations

You can back up and restore the configurations of the resource pool and host servers.

Resource pool configuration

You can utilize a XenServer host's console to back up the configuration of a resource pool. Use the following command:

xe pool-dump-database file-name=<backupfile>

This file will contain pool metadata and may be used to restore a pool configuration. Use the following command, as shown in Figure 39:

xe pool-restore-database file-name=<backupfiletorestore>

Figure 39. Restoring the resource pool configuration

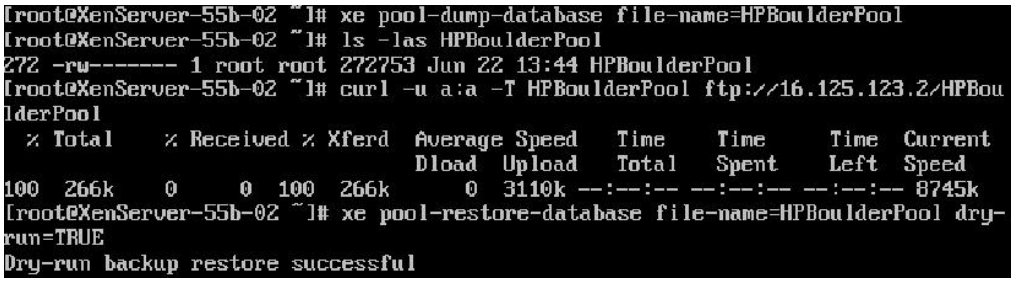
[root@XenServer-55b-02 ~]# xe pool-dump-database file-name=HPBoulderPool
[root@XenServer-55b-02 ~]# ls -las HPBoulderPool
272 -rw------- 1 root root 272753 Jun 22 13:44 HPBoulderPool
[root@XenServer-55b-02 ~]# curl -u a:a -T HPBoulderPool ftp://16.125.123.2/HPBou
lderPool
% Total % Received % Xferd Average Speed Time Time Time Current
Dload Upload Total Spent Left Speed
100 266k 0 0 100 266k 0 3110k --:--:-- --:--:-- --:--:-- 8745k
[root@XenServer-55b-02 ~]# xe pool-restore-database file-name=HPBoulderPool dry-
run=TRUE
Dry-run backup restore successful

In a restoration operation, the **dry-run** parameter can be used to ensure you are able to perform a restoration on the desired target.

For the restoration to be successful, the number of network interfaces and appropriately named NICs must match the resource pool at the time of backup.

The following **curl** command can be used to transfer files from a server to a File Transfer Protocol (FTP) server. The command is as follows:

curl –u <username>:<password> -T <filename>
ftp://<FTP_IP_address>/<Directory>/<filename>

Host configuration

You can utilize a XenServer host's console to back up the host configuration. Use the following command, as shown in Figure 40:

xe host-backup host=<host> file-name=<backupfile>

Figure 40. Backing up the host configuration

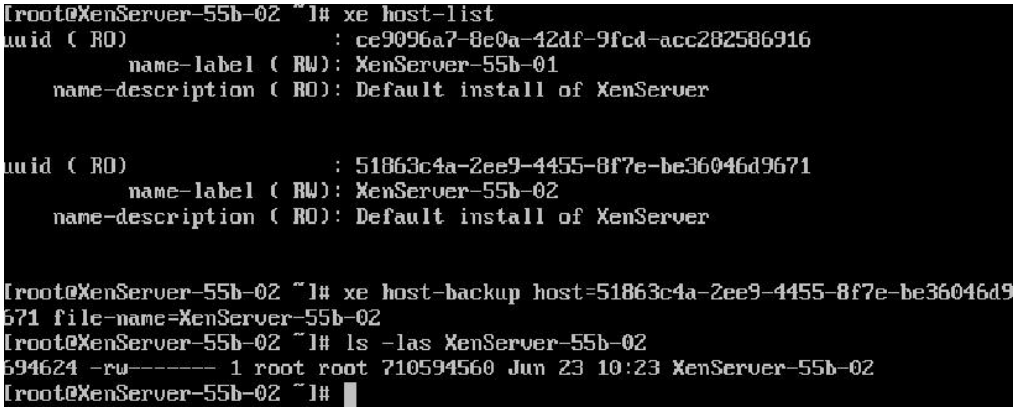
[root@XenServer-55b-02 ~]# xe host-list
uuid (RO)   : ce9096a7-8e0a-42df-9fcd-acc282586916
          name-label (RW): XenServer-55b-01
    name-description (RO): Default install of XenServer
uuid (RO)   : 51863c4a-2ee9-4455-8f7e-be36046d9671
          name-label (RW): XenServer-55b-02
    name-description (RO): Default install of XenServer
[root@XenServer-55b-02 ~]# xe host-backup host=51863c4a-2ee9-4455-8f7e-be36046d9
671 file-name=XenServer-55b-02
[root@XenServer-55b-02 ~]# ls -las XenServer-55b-02
694624 -rw------- 1 root root 710594560 Jun 23 10:23 XenServer-55b-02
[root@XenServer-55b-02 ~]#

The resulting backup file contains the host configuration and may be extremely large.

The host may be restored using the following command.

xe host-restore host=<host> file-name=<restorefile>

Original XenServer installation media may also be used for restoration purposes.

Backing up metadata

SRs contain the virtual disks used by VMs either to boot their operating systems or store data. An SR is physically connected to the hosts by physical block device (PBD) descriptors; virtual disks stored on these SRs are connected to VMs by virtual block device (VBD) descriptors.

These descriptors can be thought of as SR-level VM metadata that provide the mechanism for associating physical storage to the XenServer host and for connecting VMs to virtual disks stored on the SR. Following a disaster, the physical SRs may be available; however, you need to recreate the XenServer hosts. In this scenario, you would have to recreate the VM metadata unless this information has previously been backed up.

You can back up VM metadata using the **xsconsole** command, as shown in Figure 41. Select the desired SR.

Note
The metadata backup must be run on the master host in the resource pool, if so configured.

Figure 41. Backing up the VM metadata

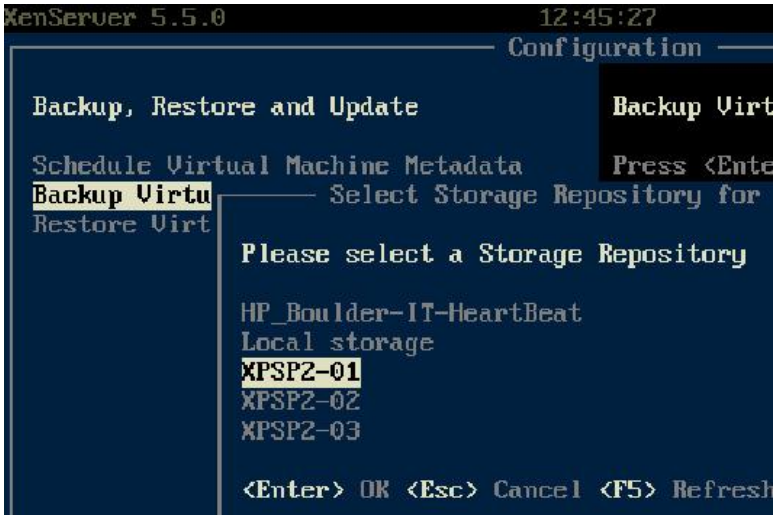

VM metadata backup data is stored on a special backup disk in this SR. The backup creates a new virtual disk image containing the resource pool database, SR metadata, VM metadata, and template metadata. This VDI is stored on the selected SR and is listed with the name **Pool Metadata Backup**.

You can create a schedule (Daily, Weekly, or Monthly) to perform this backup automatically.

The **xsconsole** command can also be used to restore VM metadata from the selected source SR. This command only restores the metadata; physical SRs and their associated data must be backed up from the storage.

SAN based Snapshots

Storage based snapshots provide storage consistent points of recovery. Initiated from the HP StorageWorks Centralized Management Controller or storage CLI, iSCSI volumes may create consistency checkpoints from which the volume may be rolled back to that point in time. The advantage of creating snapshots from the storage perspective is that very efficient features exist for working with these views of data. A snapshot can be thought of like any other volume in a point in time. The snapshot retains space efficiency, in the case of thinly provisioned snapshots, by only utilizing delta changes from writes to the original volume. In addition, snapshots and multiple levels of snapshots do not affect XenServer host efficiency to the original iSCSI volume, unlike snapshots originating from within XenServer Hosts on LVM over iSCSI volumes. Also, snapshots may be scheduled with retention schedules, for local clusters or even sent to remote locations with storage based efficiency of resource utilization. Remote snapshots have an added advantage of only pushing small deltas maximizing bandwidth availability. Remote Snapshots are discussed as the primary approach of Multi-Site Asynchronous SAN configurations.

Due to XenServer Host requirements on unique UUIDs at the storage repository and virtual disk level, the primary use of snapshots will be left to: disaster recoverability at remote sites from unique resource pools, consistency recoverability points for rolling back a volume to a point in time, and the source for creating new volume. Snapshots created from the HP StorageWorks P4000 SANs are read only consistency points that support a temporary writeable space to be mounted. This temporary white space does not persist a SAN reboot and may also be cleared manually. Since a duplicate Storage Repository must contain unique UUID SR and virtual disk data, it is impractical to manually go

thru changing this data to work with individual snapshots and at best works for only changing the original volume's UUID and persisting the old UUID with the snapshot. Best practice will suggest limiting the use of the snapshots to the previously suggested use cases. Although no storage limitation is implied with a snapshot as it is functionally equivalent to a read only volume, simplification is suggested over implementing limitless possibilities.

Recall that a Storage Repository consists of a virtual machine's virtual disk. In order to provide a consistent application state, a VM needs to be shut down or initiated in order to create a snapshot with the VSS provider. The storage volume will then be sure to have a known consistency point of data from an application and operating system perspective and will be a good candidate for initiating a storage-based snapshot, either locally to the same storage cluster or a remote snapshot. If VSS is to be relied upon for a recovery state, upon recovery, creation of a VM from the source XenCenter snapshot will be required as a recovery step.

The Storage Repository's iSCSI volume will be selected as the source for the snapshot. In this example, the VM XPSP2-05 is shut down. Highlight the XPSP2-05 volume in the CMC, right click and select New Snapshot as shown in Figure 42. The Default Snapshot Name of XPSP2-05_SS_1 will be pre-populated and by default, no servers will be assigned access. Note that if a New Remote Snapshot is selected, a Management Group will need to be selected, new remote volume name selected or created and a remote snapshot name created. It is possible to select creating a new remote snapshot and selecting the local management group thereby making a remote snapshot a local operation.

Figure 42. Select New Snapshot

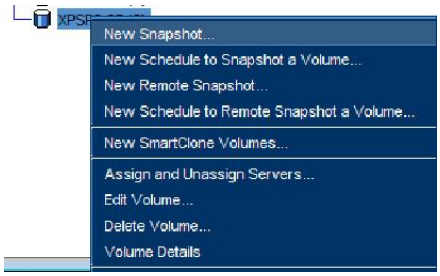

SAN based Snapshot Rollback

A storage repository that has previously been Snapshot may be rolled back to that point in time, thereby restoring the state of the virtual machine and virtual disk to its previous snapshot state. Once a determination has been made that the current virtual machine and storage is no longer valid, the storage repository upon which the virtual machine needs to recover must be reconnected to the rollback volume, as shown in Figure 43.

Figure 43. Snapshot rollback

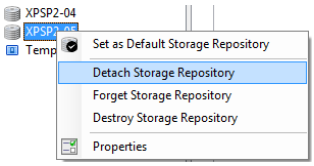

It is a best practice to disconnect from the storage repository and reattach to the new rollback storage repository; however, as long as the virtual machine is in a shut down state, the volume may simply be rolled back and virtual machine restarted to the previous state to the rolled back volume. The proper method (best practice) will be to disconnect the iSCSI session with the old volume first by first highlighting the Storage Repository in the XenCenter, right clicking and selecting **Detach Storage Repository** (see Figure 44). In the CMC, select the volume, right click and select **Rollback volume**. In XenCenter, highlight the grayed out storage repository, right click and select **Reattach Storage Repository**. Re-specify the iSCSI target portal, Discover the correct IQN appropriate for the rolled back storage repository, Discover the LUN, and select **Finish**. Ensure that **Yes** is selected to reattach the SR. In this manner, the iSCSI session is properly logged off to the storage target, the connection is broken while the storage is rolled back to the previous state, and the connection is re-established to the rolled back volume. The virtual machine may once again be restarted and will start in the state as represented by the iSCSI volume at the time of the original snapshot. See Figure 44.

Figure 44. Snapshot rollback continued

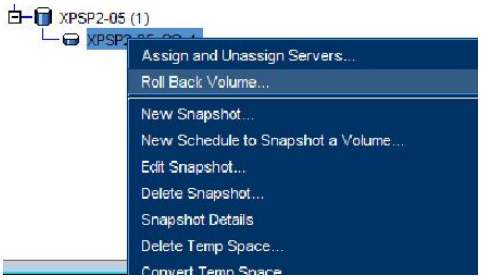

Reattach storage repositories

For new XenServer Hosts or resource pools recovering Storage Repositories, a special process exists in order to match a storage state to a virtual machine's virtual disk. In this example, an iSCSI volume contains a storage repository and virtual disk's virtual machine state, however, no configuration exists on the XenServer Host or resource pool. Select **new storage** as if attaching to a new storage repository. Of virtual disk storage type, select iSCSI type and provide the intended name. In this example, the iSCSI volume **XPSP2-05** and virtual disk has been forgotten. The name specified will be **XPSP2-05**. Provide the iSCSI target portal VIP address, discover IQNs and select the **XPSP2-05** volume, discover LUNs and finish. If a valid XenServer formatted volume exists, a warning indicating that an existing SR was found on the selected LUN. Select the option to Reattach to the storage to preserve the data currently existing on that volume. See Figure 45.

Figure 45. Reattach storage repositories

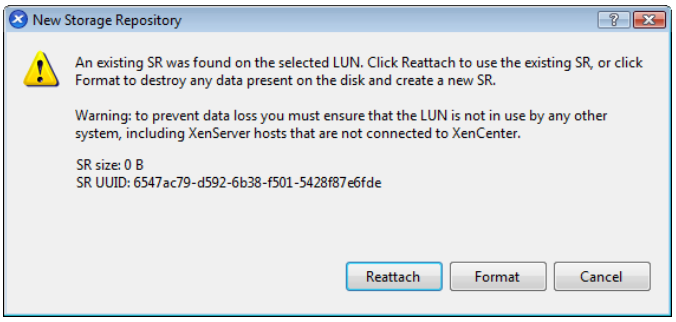
New Storage Repository
An existing SR was found on the selected LUN. Click Reattach to use the existing SR, or click Format to destroy any data present on the disk and create a new SR.
Warning: to prevent data loss you must ensure that the LUN is not in use by any other system, including XenServer hosts that are not connected to XenCenter.
SR size: 0 B
SR UUID: 6547ac79-d592-6b38-f501-5428f87e6fde
Reattach
Format
Cancel

Once the volume is reattached, a VM needs to be created of the same type and reattached to the virtual disk on that storage repository. Create a new VM, select the appropriate operating system template, provide the appropriate name. The Virtual Disks option may select anything as this will need to be manually changed. Do not select starting the VM automatically as changes still need to occur. Highlight the VM and select the **storage** tab. Detach the current virtual disk (incorrectly chosen earlier) and select **attach**. Select the **XPSP2-05 SR** and virtual disk on that SR volume. Note that (**No Name**) is the default on a reattach and may be changed to "**0**" as to XenServer defaults. The Virtual Disk name is changed on the Storage tab, properties for the XPSP2-05 SR. Select **Attach**. The VM may now be started from the state as stored on the SR. Note that SR, virtual disk, and VM uniqueness will be addressed later in this document and the requirements here specify that no cloned XPSP2-05 image exists in the resource pool or any other XenServer host seen by XenCenter.

Virtual Machines (VMs)

Creating VMs

Once storage repositories have been created for a XenServer host or resource pool, a virtual machine may be created and installed on the storage repository. It is suggested best practice to name the operating system with the same name as the storage repository as well as the iSCSI volume. In this manner, it is clear which VM is contained within which virtual disk on a storage repository stored on which iSCSI volume.

In planning for VM Storage Repository use, certain best practices must be followed for iSCSI LVM supported volumes. An LVM over iSCSI volume may contain many virtual disks; however, it is considered best practice and best performance to allocate a single storage repository to a single virtual machine. Multiple virtual disks on a single storage repository allocated to a single virtual machine is also acceptable. In a XenServer resource pool, utilizing shared storage, each XenServer host in the pool establishes an active session to the shared storage, if multiple VMs were stored on a single storage repository and each VM could be run by a different XenServer host. With HP StorageWorks P4000 clustered storage, in order to preserve volume integrity across multiple initiators to the same volume, a volume lock for writes occurs allowing only a single initiator to write at a time and thereby preserving time sequenced writes amongst multiple initiators sharing the storage. A performance bottleneck can therefore occur, which results in the recommendation for a single XenServer host to write to a single storage repository at a time. In order to enforce this best practice recommendation, a single virtual machine must be installed to a single storage repository.

Size of the storage repository

Size of the storage repository is dependent upon several factors including: size of the VM virtual disk allocated for the operating system, applications, XenCenter snapshot requirements either from the

console or from VSS enabled requestors, location of additional application data and logs (within XenServer virtual disks or separate iSCSI volumes), and planning for future growth.

An operating system installation size depends upon features chosen during the installation as well as temporary file space. Additional applications installed will also occupy space and are dependent upon what the VM applications are intended to run. Applications may also rely upon data and logging space to be available. Depending upon the architecture of a solution, separate iSCSI volumes may also be implemented for VM data stores that are mapped directly within the VM rather than externally thru the XenServer host as a storage repository. An LVM over iSCSI volume is formatted as an LVM volume virtual disk space on the LVM volume which has little overhead occupied by the LVM file system.

Snapshots require space on the original storage repository during creation. Although initially not occupying much space, changes to the original virtual disk volume over time may force the snapshot to occupy as much space as the original volume. Also, in order to utilize a snapshot, the original VM virtual disk volume space must also be available on the same storage repository to create a VM from that snapshot. In order to keep a best practice configuration, a VM created from a snapshot should then be copied to a new storage repository. This same approach will apply to VSS created snapshots. Planning for a VM must also take into consideration if snapshot features are to leave available space for snapshots or VM creations from snapshots.

Planning for future growth also minimizes the amount of administration that must occur to accommodate VM file space growth. Since HP StorageWorks iSCSI volumes can be created as Thinly Provisioned volumes, a larger storage repository than initially needed may be created and a larger virtual disk allocated to a VM than initially needed. The unused space is not carved out of the iSCSI virtualization space as the configuration is only passed down as a larger volume. By provisioning a larger than needed volume, the administrator may be able to prolong changing storage allocations and thereby save future administrative actions and thereby save time.

Increasing storage repository volume size

HP StorageWorks P4000 volumes may be edited at any time and volume size increased to accommodate future needs. In order to have an increase in volume size be recognized by XenServer hosts and ultimately VMs, several procedures must be performed. In the **HP StorageWorks CMC**, highlight the volume, right click on the volume and chose **Edit Volume**. Change the Size from the current size to the new larger size. In this example, 10GB is changed to 50GB. From the storage perspective, the volume is now larger, however, it needs to be recognized from the XenServer host. In the XenCenter console, ensure that any VM which consists of virtual disks stored on that storage repository is shut down. The storage repository must be logged out and re-logged in at the iSCSI layer to recognize the newly available storage. With the storage repository highlighted, right click on the repository and select **Detach Storage Repository**. Select **yes** to detach the repository and make the virtual disks inaccessible. Note that storage repository disk icon will become grey and the state will be detached. Viewing the iSCSI Sessions from within the CMC will note that the sessions will be logged out. In the XenCenter console with the storage repository highlight, right click and select **Reattach Storage Repository**. Keep the same name of the Storage Repository; enter the same target host of the iSCSI VIP target portal. Select **Discover IQNs** and select the same iSCSI volume that was previously detached. Select **Discover LUNs** and **Finish**. A warning will request that you ensure that you wish to reattach the SR, select **Yes**. The iSCSI volume will re-establish iSCSI sessions to all XenServer hosts in the resource pool. The reattached storage repository will now make the extra allocated space available. This space may now be used for snapshots or additional virtual disks.

VMs are allocated virtual disks. To grow a virtual disk, in the XenCenter console, highlight the newly expanded storage repository. Select the **storage** tab. Highlight the "**0**" virtual disk and select **properties**. Select the Size and Location property and increase the size allocation. In this example, a

9GB virtual disk is changed to a 20GB virtual disk. Select **OK**. The virtual disk presented to the VM will now be 20GB.

Start the VM. Depending upon the VM's operating system, different tools must be used to extend a partition and make the extra space known as a file system to the virtual machine. Different options exist, as a new partition may be created or the original partition may be expanded. Third-party tools, such as Partition Magic, also exist and may perform this function. In this example, a Windows file system' boot partition will be expanded with the PowerQuest PartitionMagic utility, as the Windows tool diskpart may only be used to expand non-system boot partitions. Note that after the VM starts with the additional virtual disk space allocated, this new space is seen in Disk Management as unallocated and un-partitioned.

In this example, a third-party utility PartitionMagic is used to resize a live Windows NTFS file system. Note the size of the expanded partition is now 20GB. Alternatively, a new partition may be created in the free space.

Figure 46 Increasing storage volume size

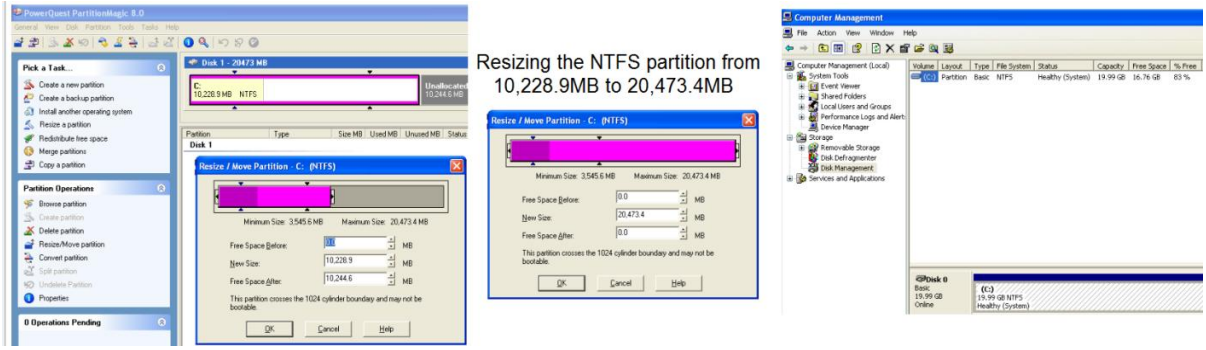

Uniqueness of VMs

Each machine on a network must be unique in order to identify one machine from another. On networks, each NIC has a unique MAC address and each a unique IP address. Within a domain, each machine a unique host name. Within Windows networks, each Windows machine caries a unique Security Identifier, SID. Within XenServer hosts or resource pools, each storage repository contains a unique UUID. With each storage repository, each virtual disk also contains a unique UUID. The purpose of each of these uniqueness attributes is to provide that instance of a machine its own identity, and a virtual machine is no different. However, virtual machines may be particularly susceptible to creating duplications as functions of replicating an entire VM and its storage become easier in virtualized environments and storage based replications in SAN environments.

With Windows-based machines, the Security Identifier is a unique name which is assigned by a Windows Domain controller during the log on process that is used to identify a subject. A SID has a format such as S-1-5-21-7623811015-3361044348-030300820-1013. Microsoft supports a mechanism for preparation of a machine for a golden image cloning process with a utility called sysprep. Note that this is the only supported way to properly clone a Windows VM. Sysprep modifies the local computer SID to make it unique to each computer. The sysprep binaries are on the Windows installation media in the \support\tools\deploy.cab file.

Process preparing a VM for Cloning

- Create, install and configure the Windows VM
- Apply Windows Updates and Service Packs
- Install Citrix XenTools paravirtualized drivers
- Install applications and any other desired customization, apply application updates
- Defragment the hard disk
- Copy the contents of the \support\tools\deploy.cab file to a \sysprep folder in the VM
- Run the sysprep utility and shutdown the VM

The virtual machine is now ready for export, copy, snapshot and create, or SAN based Snapshots or clones.

Figure 47. VM clones

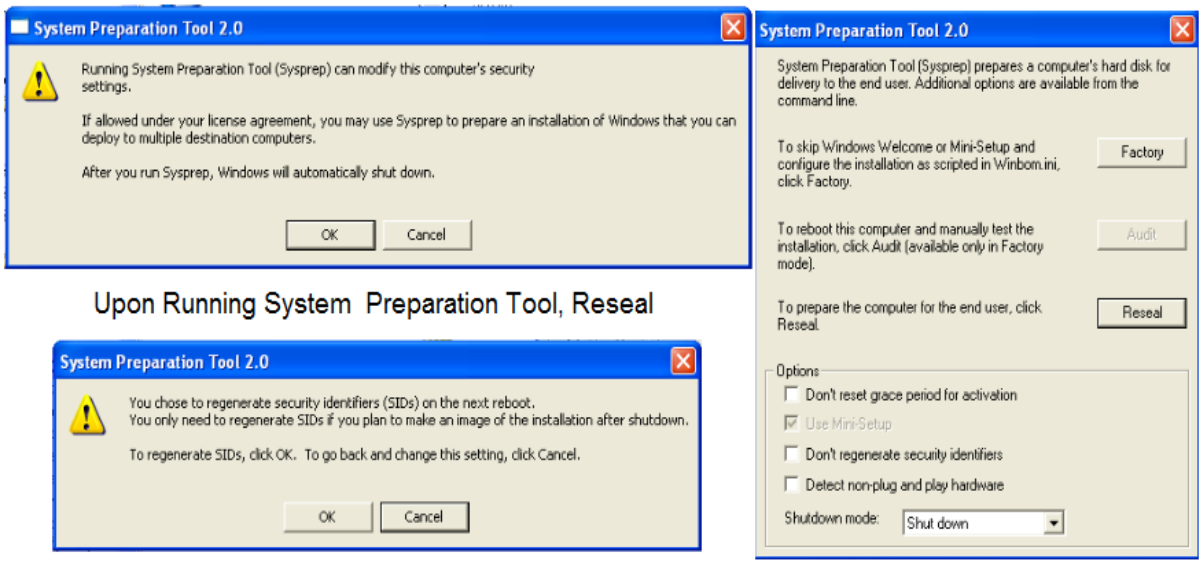
System Preparation Tool 2.0
Running System Preparation Tool (Sysprep) can modify this computer's security settings.
If allowed under your license agreement, you may use Sysprep to prepare an installation of Windows that you can deploy to multiple destination computers.
After you run Sysprep, Windows will automatically shut down.
OK
Cancel
Upon Running System Preparation Tool, Reseal
System Preparation Tool 2.0
You chose to regenerate security identifiers (SIDs) on the next reboot.
You only need to regenerate SIDs if you plan to make an image of the installation after shutdown.
To regenerate SIDs, click OK. To go back and change this setting, click Cancel.
OK
Cancel
System Preparation Tool 2.0
System Preparation Tool (Sysprep) prepares a computer's hard disk for delivery to the end user. Additional options are available from the command line.
To skip Windows Welcome or Mini-Setup and configure the installation as scripted in Winbom.ini, click Factory.
Factory
To reboot this computer and manually test the installation, click Audit (available only in Factory mode).
Audit
To prepare the computer for the end user, click Reseal.
Reseal
Options
Don't reset grace period for activation
Use Mini-Setup
Don't regenerate security identifiers
Detect non-plug and play hardware
Shutdown mode:
Shut down

Note that unsupported methods include such methods as NewSID and Symantec Ghostwalker. Each generates a unique SID and hostname on the applied image.

By utilizing XenCenter's console for exporting, copying, snapshot and creating VMs, the VM is always copied in a process that forces the new creation of Storage Repositories and virtual disks on those repositories thereby not requiring changes to their UUID.

Changing the Storage Repository and Virtual Disk UUID

A universally unique identifier, UUID, is used to identify objects inside a XenServer host or resource pool. For storage repositories, this identifier is created upon creation of a storage repository and allows one SR to be uniquely identified from another. The UUID creation algorithm ensures this uniqueness. An example UUID would be 6411ec59-7c06-3aa1-99d0-37a15a679427. A virtual disk is stored on a SR container and also contains a different UUID. This is not a problem when leveraging XenCenter data movement commands. However, HP StorageWorks P4000 SANs support many efficient features that allow instantaneous snapshots or smart clones to be created with a single backend storage command. At the storage level, one or many volumes may be created from a base volume without the data passing from SAN to XenServer host to SAN. This process is instantaneous,

may leverage space efficiency and will not tie up XenServer host resources. The downside to this process is that although a unique iSCSI volume will be created with duplicated data, the UUIDs of both the storage repository and virtual disk will also be duplicated.

Any host seen by XenCenter, including a resource pool, must not share storage repositories or virtual disks with duplicate UUIDs. This management layer depends upon uniqueness. Storage with duplicate UUIDs will not be allowed and can't be used; only the first unique UUID is seen. Therefore, a step process must be followed to force the uniqueness of the SR and virtual disk, which will allow it to be seen and used.

For this example, the system prepared XPSP2-02 VM will be used as the source. This VM is stored on a virtual disk on the XPSP2-02 storage repository. The VM is shut down. Open the HP StorageWorks Centralized Management Console (CMC), select the iSCSI volume **XPSP2-02**. Select **New Remote Snapshot** as this will create a duplicated volume copy on the SAN. Create the new Primary Snapshot. Note the default name **XPSP2-02_SS_1**. Select the Same Management Group **HP-Boulder**, select new remote volume on the existing cluster, adding a volume to the existing cluster. Select the existing **cluster IT-DataCenter** and provide a new volume name **XPSP2-02-RS-1** with appropriate replication levels (2-Way for this example). The new volume will now replicate from snapshot to snapshot and this operation occurs from within the SAN. Progress of the replicated volume may be seen on the volume's Remote Snapshots tab under the "% Complete" column. Note that this new volume XPSP2-02-RS-1 is a remote volume type and may not be used until the volume is made primary. As a remote volume type, it will be grayed out.

To use the volume, highlight the new volume **XPSP2-02-RS-1**, right click on the volume and select **edit volume**. On the Advanced tab, change the volume type from **Remot**e to **Primary**. In this example, **Thin Provisioning** is also checked. On the basic tab, select **Assign and Unassign Servers**. Ensure that all XenServer Hosts in the resource pool are assigned read/write access to the volume. In the CMC, highlight the **XPSP2-02-RS-1_RS_1** remote snapshot, right click on that snapshot and select **Delete Snapshot**. A stand-alone primary iSCSI volume, XPSP2-02-RS-1 now exists, which is a complete copy of the original XPSP2-02 volume, data, storage repository, virtual disk, and UUIDs in all.

Since only one unique UUID may be present, the choice of either forgetting the current XPSP2-02 storage repository and changing the new XPSP2-02-RS-1 repository and reattaching to the original, or changing the original XPSP2-02 repository and then re-attaching the new XPSP2-02-RS-1 must be made. In this example, detach and forget the original SR keeping its original UUID.

Step 1 - In the XenCenter Console, power down the XPSP2-02 VM. Highlight the XPSP2-02 storage repository, right click and select **Detach Storage Repository**. Select **Yes** that you want to detach this storage repository. Highlight the XPSP2-02 detached storage repository, right click and select **Forget Storage Repository**. Select **Yes** that you want to forget this storage repository. Note that the XPSP2-02 storage repository is now not listed in XenCenter.

Step 2 – In the XenCenter Console, select **New Storage**. Select the iSCSI Virtual disk storage type. Enter the iSCSI storage name XPSP2-02-RS-1, the iSCSI target portal, and select **Discover IQNs**. Select the **XPSP2-02-RS-1 iSCSI** volume. Select **Discover LUNs** and **Finish**. Select **Reattach** to preserve the existing data from the replication. Do not select format; otherwise, the VM and data on the volume will be lost. The XPSP2-02-RS-1 volume will now be attached and seen to the XenServer resource pool.

Step 3 – Open up a XenServer console in XenCenter. The XPSP2-02-RS-1 storage repository mapped device path must be found to change its UUID. On the console command line, type:

ls –las /dev/disk/by-path | grep –i "XPSP2-02-RS-1"

This command will list all the devices by path and pipe the output to grep, searching without case sensitivity for the XPSP2-02-RS-1 volume which will be found in the IQN path. In this example, the device path is /dev/disk/by-path/ip-1.1.1.225:3260-iscsi-iqn.2003-10.com.lefthandnetworks:hp-boulder:285:xpsp2-02-rs-1 as link resolving to /dev/sdd (/dev/disk/by-path/../../sdd). The

/dev/sdd is the device path that is required for the next commands and is dependent upon configuration. For example, it may be /dev/sdg or /dev/sdaa. Note the relation of the device by-path to the iSCSI IQN target name for the volume.

Step 4 – From the XenServer console in XenCenter. The XPSP2-02-RS-1 storage repository, mapped to device path /dev/sdd, is now used to locate and verify the SR UUID. Note that the appropriate device path value must be used from what was found in Step 3. On the console command line, type:

pvscan | grep –l "/dev/sdd"

The portion that is of interest is after VG_XenStorage. Highlight this value and copy it to a notepad document or write down the long UUID string. In this example: VG_XenStorage-13a7f4d6-75c7-8318-6679-eb6702b11de1.

Step 5 – From the XenServer console in XenCenter. The XPSP2-02-RS-1 storage repository physical volume attributes, mapped to device path /dev/sdd, is now changed. Note that the appropriate device path value must be used from what was found in Step 3. On the console command line, type:

pvchange --uuid /dev/sdd

The command should return that the physical volume "/dev/sdd" changed.

Step 6 – From the XenServer console in XenCenter. The XPSP2-02-RS-1 storage repository volume group attributes, mapped to the volume group VG_XenStorage-13a7f4d6-75c7-8318-6679-eb6702b11de1, is now changed. Note that the appropriate volume group path value must be used from what was found in Step 4. On the console command line, type:

vgchange --uuid VG_XenStorage-13a7f4d6-75c7-8318-6679-eb6702b11de1

The command should return that the volume group "VG_XenStorage-13a7f4d6-75c7-8318-6679-eb6702b11de1" successfully changed.

Figure 48. Volume group successfully changed

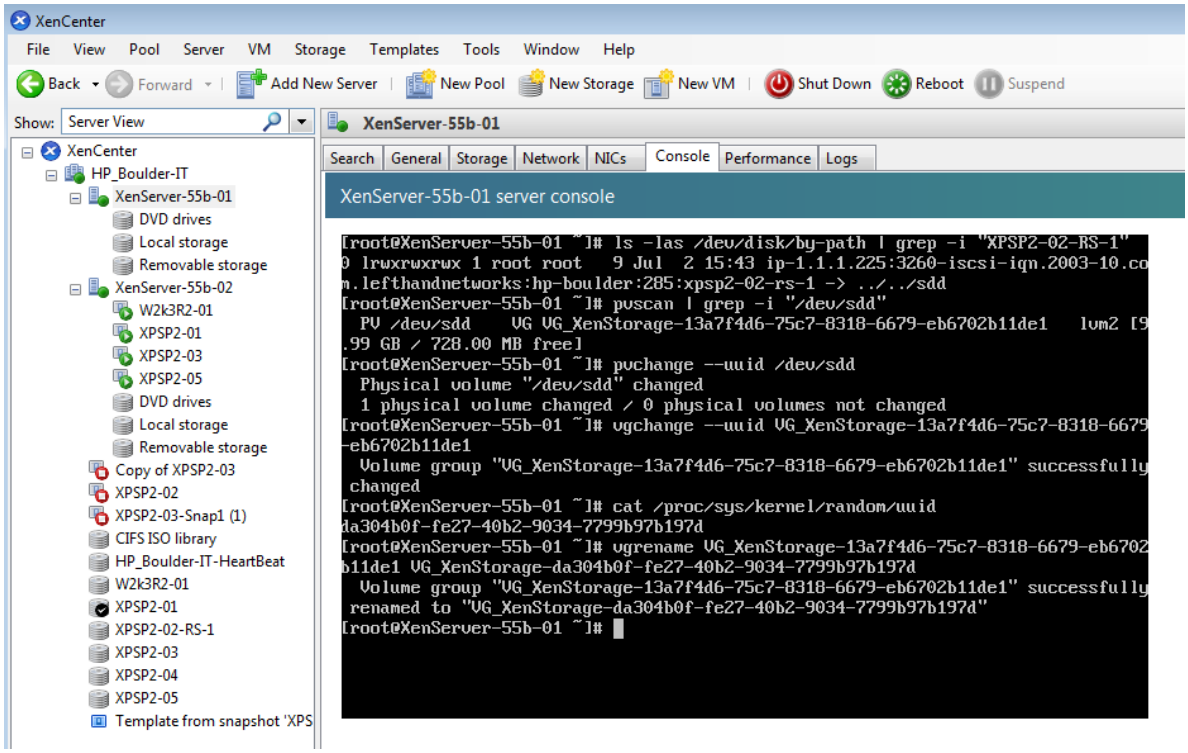

Step 7 – From the XenServer console in XenCenter. The XPSP2-02-RS-1storage repository volume group name, VG_XenStorage-13a7f4d6-75c7-8318-6679-eb6702b11de1, will be renamed to represent a new UUID for the storage repository. The "VG_XenStorage-13a7f4d6-75c7-8318-6679-eb6702b11de1" will be changed to "VG_XenStorage-13a7f4d6-75c7-8318-6679-eb6702b11de2". Note that a unique UUID may be chosen by altering a single last alphanumeric. A digit 0 thru 9 and a letter are valid characters for the UUID. Note that although naming is not enforced, it is strongly urged to keep the same number of characters. If many UUIDs are to be generated, a random UUID may be created by the following command:

cat /proc/sys/kernel/random/uuid

The command returns a random UUID. In this example, da304b0f-fe27-40b2-9034-7799b97b197d is returned and will be equally as valid to use. Either by random selection or manual choice, a unique UUID must be used. The format of the renamed command will append VG_XenStorage- to the start of the UUID. On the console command line, type:

vgrename **VG_XenStorage-13a7f4d6-75c7-8318-6679-eb6702b11de1 VG_XenStorage-da304b0f-fe27-40b2-9034-7799b97b197d**

The command returns that the volume group "VG_XenStorage-13a7f4d6-75c7-8318-6679-eb6702b11de1" is successfully renamed to "VG_XenStorage-da304b0f-fe27-40b2-9034-7799b97b197d". Note the new UUID highlighted in bold, which matches the generated UUID.

Step 8 – From the XenServer console in XenCenter. The XPSP2-02-RS-1 storage repository volume group contains a new name, "VG_XenStorage-da304b0f-fe27-40b2-9034-7799b97b197d." Now that the storage repository has been changed, the virtual disk contained on the storage repository will also need to be changed. If additional virtual disks are also contained on the same storage repository, Step 9 will need to be repeated for every logical volume found on the newly renamed storage repository volume group. On the console command line, type:

lvdisplay | grep “VG_XenStorage-da304b0f-fe27-40b2-9034-7799b97b197d”

This example will only contain two virtual disks. The command returns two names starting with VHD and concatenated with the virtual disk UUID. Note that in XenServer 5.0, the names started with LV, as shown in Figure 49.

Figure 49.New name.

Step 9 – From the XenServer console in XenCenter. The XPSP2-02-RS-1 storage repository volume group’s virtual disks need to be renamed. In this example, the two virtual disks, “/dev/VG_XenStorage-da304b0f-fe27-40b2-9034-7799b97b197d/VHD-ed07c314-5f69-491d-ba12-44f24522345a” and “/dev/VG_XenStorage-da304b0f-fe27-40b2-9034-7799b97b197d/VHD-1d128180-3ef3-4e62-977a-2d2883551058” will be renamed. Two random UUIDs will be created with the following command:

cat **/proc/sys/kernel/random/uuid; cat /proc/sys/kernel/random/uuid**

The command returns two random UUIDs. In this example, the two random UUIDs are “1a1ccad1-5528-4809-8c3c-28665474364b” and “94d23675-8e6a-460e-998a-04c0adbb47dd.” On the console command line, type each command separately:

lvrename **/dev/VG_XenStorage-da304b0f-fe27-40b2-9034-7799b97b197d/VHD-ed07c314-5f69-491d-ba12-44f24522345a /dev/VG_XenStorage-da304b0f-fe27-40b2-9034-7799b97b197d/VHD-1a1ccad1-5528-4809-8c3c-28665474364b**

lvrename /dev/VG_XenStorage-da304b0f-fe27-40b2-9034-7799b97b197d/VHD-1d128180-3ef3-4e62-977a-2d2883551058 /dev/VG_XenStorage-da304b0f-fe27-40b2-9034-7799b97b197d/VHD-94d23675-8e6a-460e-998a-04c0adbb47dd

The command returns that the each volume group has been renamed. Note the new UUID highlighted in bold which matches the generated UUID.

Figure 50. Each volume group renamed

```
[root@XenServer-55b-01 ~]# cat /proc/sys/kernel/random/uuid; cat /proc/sys/kerne
l/random/uuid
1a1ccad1-5528-4809-8c3c-28665474364b
94d23675-8e6a-460e-998a-04c0adbb47dd
[root@XenServer-55b-01 ~]# lvrename /dev/VG_XenStorage-da304b0f-fe27-40b2-9034-7
799b97b197d/VHD-ed07c314-5f69-491d-ba12-44f24522345a /dev/VG_XenStorage-da304b0f
-fe27-40b2-9034-7799b97b197d/VHD-1a1ccad1-5528-4809-8c3c-28665474364b
  Renamed "VHD-ed07c314-5f69-491d-ba12-44f24522345a" to "VHD-1a1ccad1-5528-4809-
8c3c-28665474364b" in volume group "VG_XenStorage-da304b0f-fe27-40b2-9034-7799b9
7b197d"
[root@XenServer-55b-01 ~]# lvrename /dev/VG_XenStorage-da304b0f-fe27-40b2-9034-7
799b97b197d/VHD-1d128180-3ef3-4e62-977a-2d2883551058  /dev/VG_XenStorage-da304b0
f-fe27-40b2-9034-7799b97b197d/VHD-94d23675-8e6a-460e-998a-04c0adbb47dd
  Renamed "VHD-1d128180-3ef3-4e62-977a-2d2883551058" to "VHD-94d23675-8e6a-460e-
998a-04c0adbb47dd" in volume group "VG_XenStorage-da304b0f-fe27-40b2-9034-7799b9
7b197d"
[root@XenServer-55b-01 ~]#
```

Step 10 – In XenCenter, highlight the XPSP2-02-RS-1 storage repository. Right click on the storage repository and select **Detach Storage Repository**. Select **Yes** that the storage repository is to be detached. Right click on the storage repository and select **Forge Storage Repository**. Select **Yes** that the storage repository is to be forged.

Step 11- In the XenCenter Console, select **New Storage**. Select the iSCSI Virtual disk storage type. Enter the iSCSI storage name XPSP2-02-RS-1, the iSCSI target portal, and select **Discover IQNs**. Select the **XPSP2-02-RS-1 iSCSI** volume. Select **Discover LUNs** and **Finish**. Select **Reattach** to preserve the existing data from the replication. Note the new UUID of the storage repository. Do not select format; otherwise the VM and data on the volume will be lost. The XPSP2-02-RS-1 volume will now be attached and seen to the XenServer resource pool.

Figure 51. Volume now attached to the XenServer Resource Pool

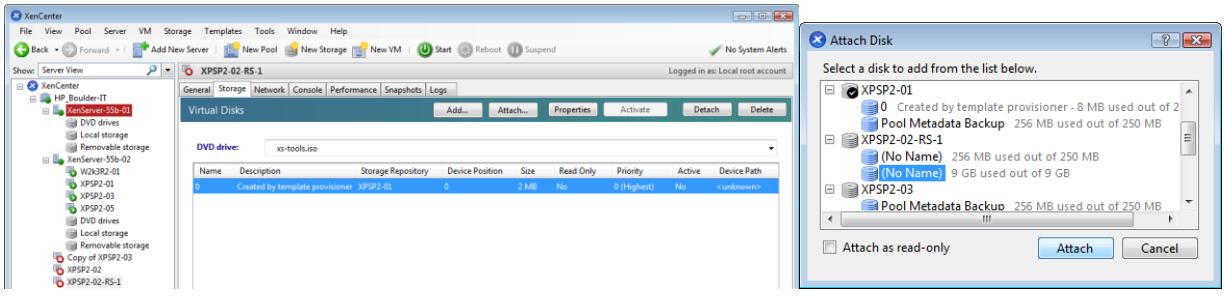

Step 12 – In the XenCenter Console, select **New VM**. Select the template appropriate for the cloned VM on the XPSP2-02-RS-1 storage repository. In this example, Windows XP SP2 is selected. Enter the name for the VM, **XPSP2-02-RS-1**. Select an ISO image. Note that an ISO image will not be required as the cloned operating system will already be installed on the virtual disk on the cloned iSCSI storage repository. Select the location for the virtual machine and the vCPUs and memory. Leave the default virtual disk as this will need to be edited later and changed to use the previously reattached

virtual disk on the XPSP2-02-RS-1 storage repository. Note that the assumption from the New VM Wizard is that a new operating system installation will be required on a new virtual disk. Select the appropriate virtual network interfaces and virtual networks. Do not start the VM automatically as the virtual disk change will need to occur first. Finish the New VM Wizard creation. Highlight the new **XPSP2-02-RS-1 VM** and select the **Storage** tab. Detach the virtual disk created by the Wizard. Select **Yes** to detach the disk. Select **Attach**. Select the **XPSP2-02-RS-1 storage repository** and select the **(No Name) 9GB** virtual disk and then select **Attach** to connect the XPSP2-02-RS-1 VM to that virtual disk. This VM is now ready to be started.

Step 13 – In the XenCenter Console, select **New Storage**. Select the **iSCSI Virtual disk** storage type. Enter the iSCSI storage name **XPSP2-02**, the iSCSI target portal, and select **Discover IQNs**. Select the **XPSP2-02 iSCSI** volume. Select **Discover LUNs** and **Finish**. Select **Reattach** to preserve the existing data from the replication. Note the new UUID of the storage repository. Do not select format; otherwise, the VM and data on the volume will be lost. The XPSP2-02 volume will now be attached and seen by the XenServer resource pool. Highlight the original **XPSP2-02 VM** and select the **Storage** tab. Select **Attach**. Select the XPSP2-02 storage repository and select the (**No Name) 9GB virtual disk** and then select **Attach** to connect the XPSP2-02 VM to that virtual disk. This VM is now ready to be started.

SmartClone the Golden Image VM

HP StorageWorks P4000 SAN SmartClone volumes are space-efficient copies of existing volumes or snapshots that are created instantaneously and are fully featured, writeable volumes. SmartCloned volumes share the initial volume space. By creating a sysprep VM and using the SmartClone process, you instantaneously create multiple volumes with access to the same base operating system configuration with only a single instance of the configuration occupying on the SAN. Each SmartClone volume persists its local writes and will only store differences from the original volume. From the HP StorageWorks P4000 Centralized Management Console (CMC), a single SmartClone operation can create up to 25 unique volumes based on the original image with the command line interface allowing more in a single command.

Figure 52. Initiating the SmartClone process

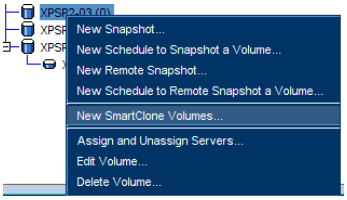

For this example, the XPSP2-03 VM is shut down. From the CMC, highlight the **XPSP2-03** volume. Right click on the highlighted iSCSI volume and select **New SmartClone Volumes**. Select **New Snapshot**. On the SmartClone Volume Setup, select a Server. Note that only a single server can be defined access and if SmartClone volumes are to be seen by multiple servers in a resource pool, the additional servers will need to be assigned access to the SmartClone volumes. Ensure that **Thin Provisioning** is selected and change the quantity to a max of **25**. For this example, a quantity of 5 will be demonstrated. Once the template is configured, select update table. Note that the SmartClone volumes will be based off the base name, VOL_XPSP2-03_SS_1_1 thru VOL_XPSP2-03_SS_1_5. Note the relationship in the CMC once created.

Figure 53.New SmartClone Volumes

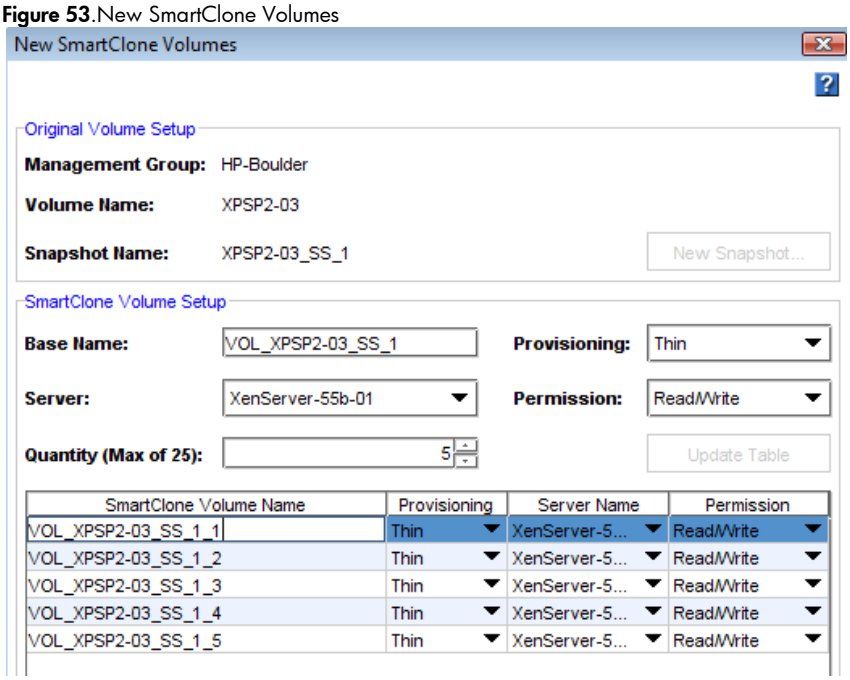

Figure 54. Five volumes

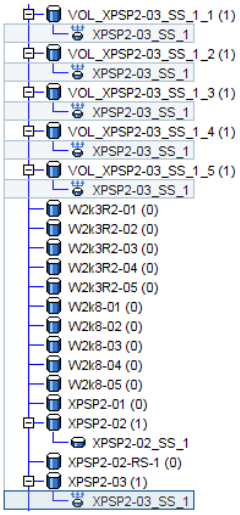

All 5 of these SmartClone volumes are unique volumes with the original single volume occupying space on the SAN. Each of these volumes may be introduced into the XenServer resource pool as identified in the earlier step. A single golden image of an operating system now serves as the source image for these 5 VMs. Modifications to the UUIDs will persist in its own volume space occupying only what is newly written in its space on the SAN. Note that each iSCSI volume is addressed thru its own IQN just like a regular volume.

Since SmartClones are based from a source snapshot, each VM is now managed as a single VM entity. If single-point-patch management is required, the original VM's volume must be patched and new SmartClone VMs must be recreated. A single base snapshot cannot be patched to roll changes into the SmartClones based upon that snapshot. This important distinction classifies SmartClone space saving and instant image creations targeted towards speeding initial deployment. Note that although

initial deployment of SmartClone volumes takes no additional footprint on the SAN, these volumes are fully writeable and may ultimately be completely re-written to occupy an entire volume's worth of space. Functions such as defragmentation at the file system level may count as additional new writes to the SAN as some operating systems prefer to write new blocks over claiming original blocks. Therefore, it is considered best practice to defragment before a SmartClone is performed. SmartClone volumes should disable defragmentation operations as this may lead to volumes filling out their thin provision.

For more information

HP StorageWorks P4000 SANs
http://h18000.www1.hp.com/products/storageworks/p4000/

HP StorageWorks P4000 Manuals
HP StorageWorks Networks Quick Start Guide
http://h20000.www2.hp.com/bizsupport/TechSupport/DocumentIndex.jsp?contentType=SupportManual&lang=en&cc=us&docIndexId=64179&taskId=101&prodTypeId=12169&prodSeriesId=3936136

Citrix XenServer
http://h71019.www7.hp.com/ActiveAnswers/cache/457122-0-0-225-121.html

http://www.citrix.com/English/ps2/products/product.asp?contentID=683148

HP Virtualization with Citrix (best practices and reference configurations)
http://www.hp.com/go/citrix

HP Solution Centers
http://www.hp.com/products1/solutioncenters/?jumpid=go/solutioncenters

To help us improve our documents, please provide feedback at http://h20219.www2.hp.com/ActiveAnswers/us/en/solutions/technical_tools_feedback.html.

Technology for better business outcomes

© Copyright 2010 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

4AA0-3524ENW, February 2010